

Asymptotic optimality of Tailored Base-Surge policies in dual-sourcing inventory systems

Linwei Xin

Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801

David A. Goldberg

School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332

Dual-sourcing inventory systems, in which one supplier is faster (i.e. express) and more costly, while the other is slower (i.e. regular) and cheaper, arise naturally in many real-world supply chains. These systems are notoriously difficult to optimize due to the complex structure of the optimal solution and the curse of dimensionality, having resisted solution for over 40 years. Recently, so-called Tailored Base-Surge (TBS) policies have been proposed as a heuristic for the dual-sourcing problem. Under such a policy, a constant order is placed at the regular source in each period, while the order placed at the express source follows a simple order-up-to rule. Numerical experiments by several authors have suggested that such policies perform well as the lead time difference between the two sources grows large, which is exactly the setting in which the curse of dimensionality leads to the problem becoming intractable. However, providing a theoretical foundation for this phenomenon has remained a major open problem.

In this paper, we provide such a theoretical foundation by proving that a simple TBS policy is indeed asymptotically optimal as the lead time of the regular source grows large, with the lead time of the express source held fixed. Our main proof technique combines novel convexity and lower-bounding arguments, an explicit implementation of the vanishing discount factor approach to analyzing infinite-horizon Markov decision processes, and ideas from the theory of random walks and queues, significantly extending the methodology and applicability of a novel framework for analyzing inventory models with large lead times recently introduced by Goldberg and co-authors in the context of lost-sales models with positive lead times.

Key words: inventory, dual-sourcing, Tailored Base-Surge policy (TBS), lead time, asymptotic optimality, convexity.

1. Introduction

A common practice in the management of global supply chains is *dual-sourcing* (cf. Rao, Scheller-Wolf and Tayur (2000)). Under a dual-sourcing strategy, the companies usually purchase their materials from a regular supplier at a lower cost, but they are also able to obtain materials from an expedited supplier at a higher cost under emergency circumstances. For example, in the summer of 2003, Amazon used FedEx to deliver the new Harry Potter more promptly and maintained regular shipping via UPS (cf. Kelleher (2003), Veeraraghavan and Scheller-Wolf (2008)). Allon and Van Mieghem (2010) describes an example of a \$10 billion high-tech U.S. company that has two suppliers, one in Mexico and one in China. The one in Mexico has shorter lead time but higher per-unit ordering cost; the one in China has longer lead time (5 to 10 times longer) but lower per-unit ordering cost. The company takes advantage of the dual-sourcing strategy to meet the demand more responsively (from Mexico) as well as less expensively (from China).

Although dual-sourcing is attractive, and very relevant to practice, optimizing a dual-sourcing inventory system is notoriously challenging. Such inventory systems have been studied now for over forty years, but the structure of the optimal policy remains poorly understood, with the exception of when the system is consecutive, i.e., the lead time difference between the two sources is exactly one. More specifically, the earliest studies of periodic review dual-sourcing inventory models include Barankin (1961), Daniel (1963), and Neuts (1964), which showed that base-stock (also known as order-up-to) policies are optimal when the lead times of the two sources are zero and one respectively. Fukuda (1964) extended the result to general lead time settings as long as the lead time difference remains one. Whittmore and Saunders (1977) showed that the optimal policy is no longer a simple base-stock policy when the lead time difference is beyond one and the structure of the optimal policy can be quite complex. Furthermore, it is well known that a dual-sourcing inventory system can be regarded as a generalization of a lost-sales inventory system (cf. Sheopuri, Janakiraman and Seshadri (2010)). Indeed, the intractability of both the dual-sourcing and lost-sales inventory models has a common source - as the lead time grows, the state-space of the

natural dynamic programming (DP) formulation grows exponentially, rendering such techniques impractical. This issue is typically referred to as the “curse of dimensionality” (cf. Karlin and Scarf (1958), Morton (1969), Zipkin (2008)), and we refer the reader to Goldberg et al. (2015) and Xin and Goldberg (2015) for a relevant discussion in the context of lost-sales inventory models.

There is a vast literature investigating periodic review dual-sourcing inventory models as well as their variants, and we refer the interested reader to the survey of Minner (2003), as well as e.g. the more recent works of Feng et al. (2006), Fox, Metters and Semple (2006), Chen, Xue and Yang (2013), Huggins and Olsen (2010), Angelus and Özer (2015), Boute and Van Mieghem (2015), Gong, Chao and Zheng (2014), Song and Zipkin (2009), and the references therein.

As an exact solution seems out of reach, the operations research and management communities have instead investigated certain structural properties of the optimal policy (cf. Hua et al. (2014)), and exerted considerable effort towards constructing various heuristic policies. Veeraraghavan and Scheller-Wolf (2008) proposed the family of dual index (DI) policies, which have two base-stock levels, one for the regular source and one for the express source, and “orders up” to bring appropriate notions of inventory position up to these levels. Scheller-Wolf, Veeraraghavan and van Houtum (2008) analyzed the closely related class of single index (SI) policies, for which the relevant notions of inventory position are different. Both families of policies seem to perform well in numerical studies. Sheopuri, Janakiraman and Seshadri (2010) considered two generalized classes of policies: one with an order-up-to structure for the express source, and one with an order-up-to structure for the regular source. Their numerical experiment showed that such policies can outperform DI policies. In the presence of production capacity costs, Boute and Van Mieghem (2015) studied dual-sourcing smoothing policies, under which the order quantities from both sources in each period are convex combinations of observed past demands. They analyzed such policies under normally distributed demand, and their numerical results showed that these policies performed better for higher capacity costs and longer lead time differences (between the two sources).

A simple and natural policy that is implemented in practice, which will be the subject of our own investigations, is the so-called *Tailored Base-Surge* (TBS) policy. It was first proposed and analyzed in Allon and Van Mieghem (2010), where we note that closely related standing order policies had been studied previously (cf. Rosenshine and Obee (1976), Janssen and De Kok (1999)). Under such a TBS policy, a constant order is placed at the regular source in each period to meet a *base* level of demand, while the orders placed at the express source follow an order-up-to rule to manage demand *surges*. We refer to Mini-Case 6 in Van Mieghem (2008) for more about the motivation and background of TBS policies. Note that dual-sourcing inventory systems in which a constant-order policy is implemented for the regular source are essentially equivalent to single-sourcing inventory systems with constant returns, which have been investigated in the literature (cf. Fleischmann and Kuik (2003), DeCroix, Song and Zipkin (2005)).

Allon and Van Mieghem (2010) analyzed TBS policies in a continuous review model, and their focus was to find the best TBS policy. Numerical results in Klosterhalfen, Kiesmüller and Minner (2011) and Rossi, Rijpkema and van der Vorst (2012) showed that TBS policies are comparable to DI policies, and outperform DI policies for some problem instances. Allon and Van Mieghem (2010) conjectured that this policy performs more effectively as the lead time difference between the two sources grows. Janakiraman, Seshadri and Sheopuri (2015) (henceforth denoted JSS) analyzed a periodic review model and studied the performance of the TBS policy. They provided an explicit bound on the performance of TBS policies compared to the optimal one when the demand had a specific structure, and provided numerical experiments suggesting that the performance of the TBS policy improves as the lead time difference grows large.

However, to date there is no theoretical justification for the good behavior of TBS policies as the lead time difference grows large, and giving a solid theoretical foundation to this observed phenomena remains a major open question. We note that until recently, a similar state of affairs existed regarding the good performance of constant-order policies as the lead time grows large in single-source lost-sales inventory models. However, using tools from applied probability, queueing theory, and convexity, this phenomena was recently explained in Goldberg et al. (2015) and

Xin and Goldberg (2015), in which it was proven that a simple constant-order policy is asymptotically optimal in this setting as the lead time of the single source grows large. The intuition here is that as the lead time grows large, so much randomness is introduced into the system between when an order is placed and when that order is received, that it is essentially impossible for any algorithm to meaningfully use the state information to make significantly better decisions. Thus a policy which ignores the state information (i.e. constant-order policy) performs nearly as well as an optimal policy. We note that the results of Xin and Goldberg (2015) further demonstrate that the optimality gap of the constant-order policy actually shrinks exponentially fast to zero as the lead time grows large, and provide explicit and effective bounds even for moderate-to-small lead times.

1.1. Our contributions

In this paper, we resolve this open question by proving that, when the lead time of the express source is held fixed, a simple TBS policy is asymptotically optimal as the lead time of the regular source grows large. Our results provide a solid theoretical foundation for the conjectures and numerical experiments of Allon and Van Mieghem (2010) and JSS. Interestingly, the simple TBS policy performs nearly optimally exactly when standard DP-based methodologies become intractable due to the aforementioned “curse of dimensionality”. Furthermore, as the “best” TBS policy can be computed by solving a convex program that does not depend on the lead time of the regular source (cf. JSS), our results lead directly to very efficient algorithms (with complexity independent of the lead time of the regular source) with asymptotically optimal performance guarantees. We also explicitly bound the optimality gap of the TBS policy for any fixed lead time (of the regular source), and prove that this decays inverse-polynomially in the lead time of the regular source. Perhaps most importantly, since many companies are already implementing such TBS policies (cf. Allon and Van Mieghem (2010)), our results provide strong theoretical support for the widespread use of TBS policies in practice. Our main proof technique combines novel convexity and lower-bounding arguments, an explicit implementation of the vanishing discount factor approach to

analyzing infinite-horizon Markov decision processes (MDP), and ideas from the theory of random walks and queues. Our methodology significantly extends the framework for analyzing inventory models with large lead times recently introduced in Goldberg et al. (2015) and Xin and Goldberg (2015) in the context of lost-sales models with positive lead times. Indeed, in the present work we relate the performance of an optimal policy to a certain *dynamic* optimization problem by applying the conditional Jensen’s inequality, while in Xin and Goldberg (2015) the relevant optimal policy could be bounded by a *static* optimization problem after applying the (non-conditional) Jensen’s inequality. The inherently dynamic nature of the resulting bounds introduce several additional difficulties not encountered previously, and which we address in the present work.

1.2. Outline of paper

The rest of the paper is organized as follows. We formally define the dual-sourcing problem in Section 2, and describe the TBS policy in Section 2.1. We state our main result in Section 2.2, and prove our main result in Section 3. We summarize our main contributions and propose directions for future research in Section 4. We also include a technical appendix in Section 5.

2. Model description, problem statement and assumptions

In this section, we formally define our dual-sourcing inventory problem, closely following the definitions given in Sheopuri, Janakiraman and Seshadri (2010). Let $\{D_t\}_{t \in (-\infty, \infty)}, \{D'_t\}_{t \in (-\infty, \infty)}$ be mutually independent sequences of nonnegative independent and identically distributed (i.i.d.) demand realizations, distributed as the non-negative random variable (r.v.) D , which we assume to have finite mean, and (to rule out certain trivial degenerate cases) to have strictly positive (possibly infinite) variance. Here we have introduced two doubly indexed sequences to prevent any possible confusion regarding dependencies of various demand realizations. Let \hat{G} be an independent geometrically distributed r.v., where $\mathbb{P}(\hat{G} = k) = 2^{-k}, k \geq 1$. As a notational convenience, let us define all empty sums to equal zero, empty products to equal one, $\frac{1}{\infty} = 0$, $\mathbf{0}(\mathbf{1})$ denote the all zeros (ones) vector, and $\mathbb{I}(A)$ denote the indicator of the event A . Let $L \geq 1$ be the deterministic lead time of the regular source (R), and $L_0 \geq 0$ the deterministic lead time of the express source (E),

where $L > L_0 + 1$. Let c_R, c_E be the unit purchase costs of the regular and express sources, and h, b be the unit holding and backorder costs respectively, with $c \triangleq c_E - c_R > 0$. In addition, let I_t denote the on-hand inventory at the start of period t (before any orders or demands are received), and $q_t^R(q_t^E)$ denote the order placed from R(E) at the beginning of period t . Note that due to the leadtimes, the order received from R(E) in period t is $q_{t-L}^R(q_{t-L_0}^E)$. As we will be primarily interested in the corresponding long-run-average problem, and for simplicity (in later proofs), we suppose that the initial conditions are such that (s.t.) the initial inventory is $-\sum_{i=1}^{\hat{G}} D'_{-i}$, and no initial orders have been placed from either R or E. Indeed, the associated system state will prove convenient to use as a “regeneration point” when analyzing certain Markov chains which arise in our proofs, where we note that the geometric distribution allows us to preclude certain kinds of pathological periodic / lattice behavior which might otherwise interfere with proving the existence of relevant stationary measures. We note that although assuming such a convenient randomized initial condition simplifies several technical proofs along these lines, such an assumption is not strictly necessary for our analysis, since the associated long-run average problem is insensitive to the particular choice of initial conditions.

As a notational convenience, we define $q_k^R = q_k^E = 0, k \leq 0$. For $t = 1, \dots, T$, the events in period t are ordered as follows.

- Ordering decisions from R and E are made (i.e. q_t^E, q_t^R are chosen);
- New inventory $q_{t-L}^R + q_{t-L_0}^E$ is delivered and added to the on-hand inventory;
- The demand D_t is realized, costs for period t are incurred, and the inventory is updated.

Note that the on-hand inventory is updated according to $I_{t+1} = I_t + q_{t-L}^R + q_{t-L_0}^E - D_t$, and may be negative since backorder is allowed.

We now formalize the family of admissible policies Π , which will determine the new orders placed. An admissible policy π consists of a sequence of measurable maps $\{f_t^\pi, t \geq 1\}$, where each f_t^π is a deterministic measurable function with domain \mathbb{R}^{L+L_0+1} and range $\mathbb{R}^{+,2}$. In that case, for a given policy π , the regular order placed in period t equals $f_{R,t}^\pi(q_{t-L}^R, \dots, q_{t-1}^R, q_{t-L_0}^E, \dots, q_{t-1}^E, I_t)$; while

the express order placed in period t equals $f_{E,t}^\pi(q_{t-L}^R, \dots, q_{t-1}^R, q_{t-L_0}^E, \dots, q_{t-1}^E, I_t)$, and Π denotes the family of all such admissible policies π .

Let $G(y)$ be the sum of the holding and backorder costs when the inventory level equals y in the end of a time period, i.e. $G(y) \triangleq hy^+ + by^-$, where $x^+ \triangleq \max(x, 0)$, $x^- \triangleq \max(-x, 0)$. Here we note that G is convex and Lipschitz, and for $x, y \in \mathbb{R}$,

$$|G(x) - G(y)| \leq \max(b, h)|x - y| \quad , \quad |G(x)| \geq \min(b, h)|x|. \quad (1)$$

For $t \geq L_0 + 1$, let C_t be the sum of the holding and backorder costs incurred in time period t , plus the ordering cost incurred for orders placed in period $t - L_0$, i.e. $C_t \triangleq c_R q_{t-L_0}^R + c_E q_{t-L_0}^E + G(I_t + q_{t-L}^R + q_{t-L_0}^E - D_t)$. We note that charging in period t for orders placed in period $t - L_0$ is a standard “accounting trick” in the inventory literature to simplify various notations (cf. Zipkin (2008a)), and for the problems considered without loss of generality (w.l.o.g.). To denote the dependence of the cost on the policy π , we use the notation C_t^π . Let $C(\pi)$ denote the long-run average cost incurred by a policy π , i.e. $C(\pi) \triangleq \limsup_{T \rightarrow \infty} \frac{\sum_{t=L_0+1}^T \mathbb{E}[C_t^\pi]}{T}$, where we again note that starting the relevant sum at $t = L_0 + 1$ (as opposed to $t = 1$) is w.l.o.g. for the problems considered. The value of the corresponding long-run average cost dual-sourcing inventory optimization problem is denoted by $\text{OPT}(L) \triangleq \inf_{\pi \in \Pi} C(\pi)$.

Before proceeding, it will be useful to apply certain well-known reductions to the problem at hand, where we note that similar reductions are known to hold for many classical inventory problems with backlogging (cf. Karlin and Scarf (1958), Scarf (1960)). First, as stated in cf. Sheopuri, Janakiraman and Seshadri (2010), for the long-run average cost problems which will be the focus of our analysis, any problem with $c_R > 0$ can be transformed into an equivalent problem with $c_R = 0$. As such we assume throughout that $c_R = 0$. Let us define the so-called *expedited inventory position* at time $t \geq 1$ as $\hat{I}_t \triangleq I_t + \sum_{k=t-L_0}^{t-1} q_k^E + \sum_{k=t-L}^{t-L+L_0} q_k^R$, which corresponds to the net inventory at the start of period t plus all orders to be received in periods $t, \dots, t + L_0$ (which were placed before period t), and the *truncated regular pipeline* at time t as the $(L - L_0 - 1)$ -dimensional vector $\mathcal{R}^t \triangleq (q_{t-L+L_0+1}^R, \dots, q_{t-1}^R)$, with $\mathcal{R}_k^t = q_{t-L+L_0+k}^R, k = 1, \dots, L - L_0 - 1$. Let $\hat{\Pi}$

denote those policies belonging to π with the additional restriction that the new orders q_t^R, q_t^E are measurable functions of only \hat{I}_t, \mathcal{R}^t . More formally, $\pi \in \hat{\Pi}$ if there exists a sequence of measurable maps $\{\hat{f}_t^\pi, t \geq 1\}$, where each \hat{f}_t^π is a deterministic measurable function with domain \mathbb{R}^{L-L_0} and range $\mathbb{R}^{+,2}$, s.t. the regular order placed in period t equals $\hat{f}_{R,t}^\pi(\mathcal{R}^t, \hat{I}_t)$ and the express order placed in period t equals $\hat{f}_{E,t}^\pi(\mathcal{R}^t, \hat{I}_t)$.

Note that $\hat{I}_1 = -\sum_{i=1}^{\hat{G}} D'_{-i}$, and $\mathcal{R}^1 = \mathbf{0}$. Also, for any policy $\pi \in \hat{\Pi}$ and $t \geq 1$, it holds that $\hat{I}_{t+1} = \hat{I}_t + q_t^E + \mathcal{R}_1^t - D_t$, $\mathcal{R}_k^{t+1} = \mathcal{R}_{k+1}^t$ for $k \in [1, L - L_0 - 2]$, and $\mathcal{R}_{L-L_0-1}^{t+1} = q_t^R$. Furthermore, for all $t \geq L_0 + 1$, $C_t^\pi = G(\hat{I}_{t-L_0} + q_{t-L_0}^E - \sum_{i=t-L_0}^t D_i) + cq_{t-L_0}^E$. Then the following is proven in Sheopuri, Janakiraman and Seshadri (2010).

LEMMA 1 (Sheopuri, Janakiraman and Seshadri (2010) Lemma 2.1). $\inf_{\pi \in \Pi} C(\pi) = \inf_{\pi \in \hat{\Pi}} C(\pi)$, i.e. one may w.l.o.g. restrict oneself to policies belonging to $\hat{\Pi}$.

For the remainder of the paper, we thus consider the relevant optimization only over policies belonging to $\hat{\Pi}$, i.e.

$$\text{OPT}(L) = \inf_{\pi \in \hat{\Pi}} C(\pi). \quad (2)$$

For a given policy $\pi \in \hat{\Pi}$, let $\mathcal{R}^{\pi,t}(\hat{I}_t^\pi)$ denote a r.v. distributed as the truncated regular pipeline (expedited inventory position) at the start of period t under policy π . Similarly, let $q_t^{\pi,E}(q_t^{\pi,R})$ denote the expedited (regular) order placed in period t , and suppose that all these r.v.s are constructed on a common probability space, and have the appropriate joint distribution induced by the operation of π over time.

2.1. TBS policy

In this section, we formally introduce the family of TBS policies, and characterize the “best” TBS policy. A TBS policy $\pi_{r,S}$ with parameters (r, S) is defined (cf. JSS) as the policy that places a constant order r from R in every period, and follows an order-up-to rule from E which in each period raises the expedited inventory position to S (if it is below S), and otherwise orders nothing. More formally, under this policy $q_t^R = r$, and $q_t^E = \max(0, S - \hat{I}_t)$, for all t .

Let $I_\infty^r \triangleq \sup_{j \geq 0} \left(jr - \sum_{i=1}^j D_i \right)$. In that case, it follows from the results of JSS that

$$C(\pi_{r,S}) = c(\mathbb{E}[D] - r) + \mathbb{E} \left[G \left(I_\infty^r + S - \sum_{i=1}^{L_0+1} D'_i \right) \right]. \quad (3)$$

Note that for each r , the minimization problem $\inf_{S \in \mathbb{R}} C(\pi_{r,S})$ is equivalent to a standard one-period newsvendor problem. Furthermore, defining $F^\infty(r) \triangleq \inf_{S \in \mathbb{R}} C(\pi_{r,S})$, it is proven in JSS that $F^\infty(r)$ is convex in r on $(-\infty, \mathbb{E}[D])$. Combining the above with standard results for single-server queues (cf. Asmussen (2003)) and (1), we conclude that there exists at least one pair (r^*, S^*) s.t. $r^* \in \arg \min_{0 \leq r \leq \mathbb{E}[D]} F^\infty(r)$ and $S^* \in \arg \min_{S \in \mathbb{R}} C(\pi_{r^*,S})$; that this pair defines the TBS policy with least long-run-average cost; and that this pair can be computed efficiently by solving a convex program which is independent of the larger lead time L .

2.2. Main result

2.3. Additional definitions and notations

Before stating our main result, we will need several additional definitions and notations to describe various relevant quantities which will appear in our bounds on the optimality gap. For $\theta \geq 0$ and $\epsilon \in (0, \mathbb{E}[D]]$, let us define

$$\phi_\epsilon(\theta) \triangleq \exp(\theta(\mathbb{E}[D] - \epsilon)) \mathbb{E}[\exp(-\theta D)] \quad , \quad \gamma_\epsilon \triangleq \inf_{\theta \geq 0} \phi_\epsilon(\theta),$$

and $\vartheta_\epsilon \in \arg \min_{\theta \geq 0} \phi_\epsilon(\theta)$ denote the supremum of the set of minimizers of $\phi_\epsilon(\theta)$, where we define ϑ_ϵ to equal ∞ if the above infimum is not actually attained. Note that $\phi_\epsilon(\theta)$ is a continuous and convex function of θ on $(0, \infty)$, and right-continuous function of θ at 0. In addition, it follows from Folland (1999) Theorem 2.27 that $\phi_\epsilon(\theta)$ is right-differentiable at zero, with derivative equal to $-\epsilon$. We conclude from the definition of derivative and a straightforward contradiction argument that $\vartheta_\epsilon > 0$ and $\gamma_\epsilon \in [0, 1)$. Let $g \triangleq \inf_{x \in \mathbb{R}} \mathbb{E} \left[G \left(x - \sum_{i=1}^{L_0+1} D'_i \right) \right] > 0$, and $U \triangleq C(\pi_{0,0}) = c\mathbb{E}[D] + \mathbb{E}[G(-\sum_{i=1}^{L_0+1} D'_i)]$, in which case it is easily verified that $g \leq \text{OPT}(L) \leq U$ for all $L > L_0 + 1$. We also make the following additional definitions:

$$p_0 \triangleq \mathbb{P}(D < \mathbb{E}[D]) \in (0, 1) \quad , \quad \hat{p}_0 \triangleq \left(\frac{1}{2} p_0 (1 - p_0) \right)^{\frac{1}{2}} \in (0, 1) \quad , \quad Q_0 \triangleq \inf \{ x \in \mathbb{R}^+ : \mathbb{P}(D \leq x) \geq \frac{1}{2} p_0 \} \in [0, \mathbb{E}[D]],$$

$$\begin{aligned}
\eta_0 &\triangleq \inf_{z \in \mathbb{R}} \mathbb{E}[|z - D|] > 0 \quad , \quad c_0 \triangleq \frac{1}{240} \min(b, h) \hat{p}_0 \eta_0 \quad , \quad U_0 \triangleq 64(L_0 + 1) \frac{\max^2(b, h)}{\min(b, h)} \mathbb{E}[D], \\
\epsilon_0 &\triangleq \min \left(\mathbb{E}[D] - Q_0, \frac{1}{4}(\eta_0 \hat{p}_0)^2, 1 - 2^{-\frac{\hat{p}_0^2}{400}}, \frac{1}{625} c_0^2 (U_0 2^{L_0} + \eta_0 + U + 1)^{-2} \right) \in (0, 1 - 2^{-\frac{1}{400}}) \subset (0, .002), \\
Y_0 &\triangleq 25g^{-2} (U_0 2^{L_0} + \max(b, h) \gamma_{\epsilon_0} \vartheta_{\epsilon_0}^{-1} (1 - \gamma_{\epsilon_0})^{-2})^2 + L_0 + 1.
\end{aligned}$$

Our main result proves that the best TBS policy is asymptotically optimal as $L \rightarrow \infty$, and provides explicit bounds on the optimality gap.

THEOREM 1. *For all $L_0 \geq 0$, $\epsilon \in (0, 1)$, and $L > \epsilon_0^{-2} + Y_0 \epsilon^{-2}$, it holds that $\frac{C(\pi_{r^*, S^*})}{OPT(L)} < 1 + \epsilon$.*

COROLLARY 1. $\lim_{L \rightarrow \infty} \frac{C(\pi_{r^*, S^*})}{OPT(L)} = 1$.

3. Proof of Theorem 1

3.1. Lower bound for the optimal cost

In this section, we prove a lower bound for $OPT(L)$ by extending the steady-state/convexity approach of Xin and Goldberg (2015) to the dual-sourcing setting. We note that here our lower bound will involve a non-trivial optimization over measurable functions, in contrast to the bounds used in Xin and Goldberg (2015) which were of a static nature. As in Xin and Goldberg (2015), we will proceed by relating the “long-run behavior” of “an optimal policy” to a certain TBS policy. At a high level, we will combine convexity and the conditional Jensen’s inequality with the fact that the r.v.s corresponding to (appropriately defined stationary versions of) the different components of the truncated regular pipeline vector (under the optimal policy) have the same mean, which will (approximately) coincide with the constant order from R in our TBS policy. Furthermore, when we apply the conditional Jensen’s inequality to certain terms corresponding to (appropriately defined stationary versions of) the expedited orders under the same optimal policy, the resulting terms will be suitably measurable functions of past demands, which will (approximately) coincide with the amount of inventory ordered from E in our TBS policy.

3.1.1. Connecting to a stationary problem. As in Xin and Goldberg (2015), our program immediately encounters a technical problem. Namely, the natural way to analyze the “long-run behavior” of an optimal policy is through the steady-state distribution of the Markov chain

induced by this policy. However, it is not obvious that this steady-state exists. Actually, it is not even obvious that there exists a stationary optimal policy (so that the dynamics even define a Markov chain), nor even that there even exists an optimal policy at all (as opposed to it only being approached). Although such questions have been rigorously analyzed for simpler inventory models in Huh, Janakiraman and Nagarajan (2011), such questions have not been rigorously answered for the setting of more complicated dual-sourcing models. We note that although in Sheopuri, Janakiraman and Seshadri (2010) it is stated in passing that many of the same results should extend to the dual-sourcing setting, no proofs are provided, and the explicit assumptions needed for such a transference are not clarified. A similarly terse exposition on related questions is provided in Hua et al. (2014). Furthermore, in none of these works is the question of existence of and convergence to relevant stationary measures discussed. To overcome this, as in Xin and Goldberg (2015), we first observe that we will not actually need a random vector which is truly the steady-state of the aforementioned Markov chain (which in principle may not exist), but only need to demonstrate the existence of a random vector which has several properties that we would want such a steady-state (if it existed) to have. We now show the existence of such a random vector. We note that although closely related questions have been studied in the MDP literature (cf. Arapostathis et al. (1993)), and perturbative approaches similar to the approach we take in our own proof are in general well-known (cf. Filar (2007)), to the best of our knowledge the desired result does not follow directly from any results appearing in the literature. As such, we include a proof for completeness in the technical appendix Section 5. We note that here the relevant analysis is considerably more challenging than that given in Xin and Goldberg (2015), due to the fact that in the dual-sourcing setting the inventory level is unbounded from below, and the associated ordering levels are not known to be uniformly bounded (in contrast to the setting considered in Xin and Goldberg (2015) for which such bounds were already proven in Zipkin (2008a)). Furthermore, although several bounds exist in the dual-sourcing literature relating order levels under an optimal policy to the inventory level at the time of ordering (cf.

Sheopuri, Janakiraman and Seshadri (2010), Hua et al. (2014)), it seems that due to the inventory being unbounded below none of those bounds are suitable for our purposes. It is also worth noting that our approach is able to side-step many of the complexities and additional assumptions (e.g. finite second moment or bounded support) often required when analyzing inventory models which are unbounded from below.

We defer all relevant proofs to the technical appendix Section 5. For two r.v.s X, Y , let $X \sim Y$ denote equivalence in distribution. Before stating our result, for the sake of building intuition, we first describe what the various r.v.s appearing in our result would correspond to “if we were to assume” (which we do not, i.e. it is not an assumption of our main results) that there exists an optimal policy which is stationary, and whose corresponding Markov chain converges to a steady-state distribution, i.e. the truncated regular pipeline and expedited inventory position converge in distribution under the operation of this optimal stationary policy. In that case, our theorem contains an $(L - L_0 - 1)$ -dimensional random vector $\chi^{*,L}$, an $(L - L_0)$ -dimensional random vector $\mathbf{q}^{*,L}$, and a r.v. $\mathcal{I}^{*,L}$, which may be interpreted as follows. Suppose one has been operating under this stationary optimal policy for a long time, say up to some very large time T , at which time the system is essentially in steady-state (again we note that this discussion is purely for the sake of building intuition, and our main results do not actually assume this). Then $\chi^{*,L}$ corresponds to the steady-state truncated regular pipeline vector under this optimal policy (at time T), i.e. $\chi_i^{*,L}$ is the regular order which enters the expedited inventory position in period $T + i - 1$. $\mathbf{q}^{*,L}$ corresponds to the steady-state vector of expedited orders to be placed over the next $L - L_0$ periods under this optimal policy, i.e. $\mathbf{q}_i^{*,L}$ is the expedited order which enters the expedited inventory position in period $T + i - 1$. Finally, $\mathcal{I}^{*,L}$ corresponds to the steady-state expedited inventory position under this optimal policy (at time T).

THEOREM 2. *For all $L_0 \geq 0$ and $L > L_0 + 1$, one may construct an $L - L_0 - 1$ -dimensional random vector $\chi^{*,L}$, an $L - L_0$ -dimensional random vector $\mathbf{q}^{*,L}$, and a random variable $\mathcal{I}^{*,L}$, as well as $\{D_i, i \geq 1\}$, on a common probability space s.t. the following are true.*

(i) *W.p.1* $(\chi^{*,L}, \mathbf{q}^{*,L})$ is non-negative. Also, $(\chi^{*,L}, \mathcal{I}^{*,L})$ is independent of $\{D_i, i \geq 1\}$, and $q_i^{*,L}$ is independent of $\{D_j, j \geq i\}$ for $i \in [1, L - L_0]$.

(ii) $\chi_i^{*,L} \sim \chi_1^{*,L}$ for $i \in [1, L - L_0 - 1]$, and $q_i^{*,L} \sim q_1^{*,L}$ for $i \in [1, L - L_0]$.

(iii) For all $k \in [1, L - L_0]$,

$$\mathcal{I}^{*,L} + \sum_{i=1}^{k-1} (q_i^{*,L} + \chi_i^{*,L} - D_i) + q_k^{*,L} - \sum_{i=k}^{k+L_0} D_i \sim \mathcal{I}^{*,L} + q_1^{*,L} - \sum_{i=1}^{L_0+1} D_i.$$

(iv) $(\chi^{*,L}, \mathbf{q}^{*,L}, \mathcal{I}^{*,L})$ has finite mean.

(v) $\mathbb{E}[\chi_1^{*,L}] + \mathbb{E}[q_1^{*,L}] = \mathbb{E}[D]$.

(vi)

$$\text{OPT}(L) \geq c(\mathbb{E}[D] - \mathbb{E}[\chi_1^{*,L}]) + \mathbb{E} \left[G \left(\mathcal{I}^{*,L} + q_1^{*,L} - \sum_{i=1}^{L_0+1} D_i \right) \right].$$

3.1.2. Vanishing discount factor approach. Although Theorem 2.(vi) relates $\text{OPT}(L)$ to a certain expectation, this expectation (as written) is not immediately amenable to analysis. To remedy this, we introduce a discount factor α to implement the so-called “vanishing discount factor” approach to analyzing infinite-horizon MDP (cf. Huh, Janakiraman and Nagarajan (2011)), which will allow for a simpler analysis when we pass to the limit as $L \rightarrow \infty$. Indeed, this discount factor will help us to analyze the lower bound which arises when we apply the conditional Jensen’s inequality, as this lower bound will itself involve the solution to a non-trivial multi-stage dynamic optimization problem. We note that the lower bound which arose when related techniques were applied to single-sourcing systems with lost sales in Xin and Goldberg (2015) only involved a static optimization problem, and thus no such discount factor was introduced. In particular, Theorem 2 immediately implies the following corollary. Let $r_L \triangleq \mathbb{E}[\chi_1^{*,L}]$.

COROLLARY 2. For all $L_0 \geq 0, L > L_0 + 1$, and $\alpha \in (0, 1)$,

$$\begin{aligned} \text{OPT}(L) &\geq c(\mathbb{E}[D] - r_L) + \frac{1-\alpha}{1-\alpha^L} \sum_{k=1}^L \alpha^{k-1} \mathbb{E} \left[G \left(\mathcal{I}^{*,L} + q_1^{*,L} - \sum_{i=1}^{L_0+1} D_i \right) \right] \\ &\geq c(\mathbb{E}[D] - r_L) + (1-\alpha) \sum_{k=1}^{L-L_0} \alpha^{k-1} \mathbb{E} \left[G \left(\mathcal{I}^{*,L} + \sum_{i=1}^{k-1} (q_i^{*,L} + \chi_i^{*,L} - D_i) + q_k^{*,L} - \sum_{i=k}^{k+L_0} D_i \right) \right]. \end{aligned}$$

3.1.3. Applying the conditional Jensen's inequality and relating to a single-source inventory model. We now apply the conditional Jensen's inequality to Corollary 2, which will allow us to lower-bound $\text{OPT}(L)$ by the optimal value of a certain finite-horizon single-source inventory model with backlogged demand. We will then relate this finite-horizon problem to an associated infinite-horizon problem, which has an optimal stationary policy. Furthermore, we will connect the behavior of such an optimal stationary policy to the performance of an associated TBS policy, ultimately allowing us to prove our main results. In particular, it follows from Theorem 2 and the independence structure of the relevant r.v.s that for $k \in [1, L - L_0]$,

$$\mathbb{E} \left[\mathcal{I}^{*,L} + \sum_{i=1}^{k-1} (q_i^{*,L} + \chi_i^{*,L} - D_i) + q_k^{*,L} - \sum_{i=k}^{k+L_0} D_i \middle| D_{[k+L_0]} \right]$$

equals

$$\mathbb{E}[\mathcal{I}^{*,L}] + \sum_{i=1}^{k-1} (\mathbb{E}[q_i^{*,L} | D_{[i-1]}] + r_L - D_i) + \mathbb{E}[q_k^{*,L} | D_{[k-1]}] - \sum_{i=k}^{k+L_0} D_i.$$

Further combining with Corollary 2, the convexity of G , and Jensen's inequality for conditional expectations, we obtain the following result.

PROPOSITION 1. *For any $\alpha \in (0, 1)$ and $L > L_0 + 1$, $\text{OPT}(L) - c(\mathbb{E}[D] - r_L)$ is at least*

$$(1 - \alpha) \sum_{k=1}^{L-L_0} \alpha^{k-1} \mathbb{E} \left[G \left(\mathbb{E}[\mathcal{I}^{*,L}] - (L_0 + 1)r_L + \sum_{i=1}^{k-1} (\mathbb{E}[q_i^{*,L} | D_{[i-1]}] - (D_i - r_L)) \right. \right. \\ \left. \left. + \mathbb{E}[q_k^{*,L} | D_{[k-1]}] - \sum_{i=k}^{k+L_0} (D_i - r_L) \right) \right]. \quad (4)$$

Note that (4) is the discounted cost incurred (during periods $L_0 + 1, \dots, L$) by the policy ordering $\mathbb{E}[q_i^{*,L} | D_{[i-1]}]$ in period i , of a single-sourcing L -period backlog inventory problem with unit holding cost h , backorder cost b , zero ordering cost, discount factor α , i.i.d. demand distributed as $D - r_L$ (which we note can be positive or negative), lead time L_0 , and initial inventory position (initial net inventory plus all entries of the initial pipeline vector) $\mathbb{E}[\mathcal{I}^{*,L}] - (L_0 + 1)r_L$ (cf. Karlin and Scarf (1958)), multiplied by $(1 - \alpha)$. Such models, and their optimal policies, have been studied in-depth in the literature (cf. Karlin and Scarf (1958), Zipkin (2000), Fleischmann and Kuik (2003)), and are well-understood (especially for the case of non-negative demand, cf. Zipkin (2000)). Let $\bar{\Pi}$

denote the family of all feasible non-anticipative policies for the aforementioned inventory problem (as it is typically defined, cf. Zipkin (2000)). For $\pi \in \bar{\Pi}$, initial inventory position $x \in \mathbb{R}$, $r \in \mathbb{R}$, and $i \geq 1$, let $C_i^\pi(r, x)$ denote the cost incurred by policy π in the aforementioned inventory problem in period $i + L_0$, if the demand in each period is i.i.d. distributed as $D - r$ (with the leadtime L_0 and costs b, h as above). For $x \in \mathbb{R}, r \in \mathbb{R}, \alpha \in (0, 1), n \geq 1$, let us define

$$V_\alpha^n(r, x) \triangleq \inf_{\pi \in \bar{\Pi}} \mathbb{E} \left[\sum_{i=1}^n \alpha^{i-1} C_i^\pi(r, x) \right]; \quad (5)$$

and

$$V_\alpha^\infty(r, x) \triangleq \inf_{\pi \in \bar{\Pi}} \mathbb{E} \left[\sum_{i=1}^{\infty} \alpha^{i-1} C_i^\pi(r, x) \right]. \quad (6)$$

As a notational convenience, we define $V_\alpha^0(r, x) = 0$, $V_\alpha^n(r, -\infty) \triangleq \inf_{x \in \mathbb{R}} V_\alpha^n(r, x)$, $V_\alpha^\infty(r, -\infty) \triangleq \inf_{x \in \mathbb{R}} V_\alpha^\infty(r, x)$. Then combining the above, we derive the following lower bound for $\text{OPT}(L)$.

LEMMA 2. For all $L_0 \geq 0, L > L_0 + 1$, and $\alpha \in (0, 1)$,

$$\text{OPT}(L) \geq c(\mathbb{E}[D] - r_L) + (1 - \alpha)V_\alpha^{L-L_0}(r_L, -\infty). \quad (7)$$

3.1.4. Overview of remainder of the proof of our main results. The remainder of the proof involves a careful analysis of the right-hand-side (r.h.s.) of (7) as $L \rightarrow \infty$, and we now sketch an outline of our approach. First, we will prove that if r_L is bounded away from $\mathbb{E}[D]$, then $V_\alpha^\infty(r_L, -\infty) - V_\alpha^{L-L_0}(r_L, -\infty)$ can be suitably bounded by a function of L which converges to 0 as $L \rightarrow \infty$. We then observe that the infinite-horizon problem associated with $V_\alpha^\infty(r_L, -\infty)$ has an optimal policy which is stationary, Markov, and of order-up-to type. Furthermore, the stochastic process induced by this optimal policy will be equivalent to that induced by a corresponding TBS policy, but possibly initialized *not according to the stationary distribution of the associated inventory process*. Then we prove that r_L is indeed bounded away from $\mathbb{E}[D]$, since otherwise we can use the theory of random walks to derive a contradiction (as $\text{OPT}(L)$ would be strictly greater than U). Finally, we combine these facts to bound various error terms under a suitable choice of α (which converges to 1 as $L \rightarrow \infty$), including a term resulting from the difference between the performance of the same TBS policy under different initializations, to prove our main result Theorem 1.

3.2. Proof of Theorem 1

We now complete the proof of Theorem 1 by formalizing the argument sketched at the end of Section 3.1. Such arguments are standard in the literature on MDP and infinite-horizon inventory control problems (cf. Iglehart (1963), Sennott (1989), Schäl (1993), Fleischmann and Kuik (2003), Feinberg (2011), Huh, Janakiraman and Nagarajan (2011)). We note that the somewhat non-standard aspect here is that the demand in each period is distributed as $D - r_L$, and thus *may be negative*. As such, the original arguments typically used to analyze the relevant quantities and prove related interchange-of-limits results (cf. Iglehart (1963)) do not directly apply. The possibility of negative demand also makes the verification of the conditions of general theorems which validate such bounds and interchange-of-limits (cf. Sennott (1989), Schäl (1993)) somewhat involved, even when these theorems are customized to the inventory setting (cf. Parker and Kapuscinski (2004), Huh, Janakiraman and Nagarajan (2011)). We note that the verification of closely related results have arisen recently in the context of analyzing inventory systems with returns, which reduce to standard inventory systems where demand can be positive or negative (cf. Fleischmann and Kuik (2003)). However, those results (which verify the technical conditions of Sennott (1989)) do not seem to extend immediately to our case, and further seem to require that the demand and ordering quantities take integer values. In light of the above, and for the sake of clarity and completeness, we now provide a self-contained proof of all necessary bounds, which (combined with Lemma 2) will complete the proof of our main result Theorem 1. We defer most proofs to the technical appendix Section 5.

We begin by stating some well-known properties of $V_\alpha^n(r, x)$ and $V_\alpha^\infty(r, x)$, which follow from the results of JSS, Karlin and Scarf (1958) and Scarf (1960). We note that although in some cases the proofs there are only explicitly given for the case of non-negative demand, as noted in Heyman and Sobel (1984) and Fleischmann and Kuik (2003), the arguments carry over to the general case (in which demand may be negative) with only trivial modification.

LEMMA 3 (JSS, Scarf (1960)). *For all $\alpha \in (0, 1)$, $r, x \in \mathbb{R}$, and $n \geq 1$,*

$$V_\alpha^n(r, x) = \inf_{y \geq x} \left(\mathbb{E} \left[G\left(y - \sum_{k=n}^{L_0+n} (D_k - r)\right) \right] + \alpha \mathbb{E} [V_\alpha^{n-1}(r, y - (D_{L_0+n} - r))] \right).$$

Furthermore, $V_\alpha^n(r, x)$ is: a convex (and thus also continuous) function of x on \mathbb{R} for each fixed n, r ; a continuous function of r on \mathbb{R} for each fixed n, x ; an increasing function of x on \mathbb{R} for each fixed n, r ; and an increasing function of n on Z^+ for each fixed x, r . In addition, the infinite-horizon problem stated in the r.h.s. of (6) admits an optimal stationary Markov policy.

Next, we bound $V_\alpha^\infty(r, x) - V_\alpha^n(r, x)$, and combine our bounds with Lemma 3 to derive some useful properties of $V_\alpha^\infty(r, x)$ and the associated optimization problem. We defer all proofs to the technical appendix Section 5. Let $\bar{S}_\alpha(r) \triangleq 4(L_0 + 1) \frac{\max(b, h)}{\min(b, h)} (|r| + \mathbb{E}[D])(1 - \alpha)^{-2}$.

LEMMA 4. For $\alpha \in (0, 1)$, $r, x \in \mathbb{R}$, and $n \geq 1$,

$$0 \leq V_\alpha^\infty(r, x) - V_\alpha^n(r, x) \leq \max(b, h) (\bar{S}_\alpha(r) + |x| + |r| + \mathbb{E}[D]) (1 + L_0 + n) (1 - \alpha)^{-2} \alpha^n, \quad (8)$$

and $V_\alpha^\infty(r, x) = \lim_{n \rightarrow \infty} V_\alpha^n(r, x)$. Furthermore, for $\alpha \in (0, 1)$ and $r \in \mathbb{R}$, $V_\alpha^\infty(r, x)$ is a finite-valued, convex, and non-decreasing function of x on \mathbb{R} . Letting $S_\alpha^\infty(r)$ denote the supremum of the set of minimizers (in x) of $V_\alpha^\infty(r, x)$, it holds that $|S_\alpha^\infty(r)| \leq \bar{S}_\alpha(r)$, and the infinite-horizon problem stated in the r.h.s. of (6) admits an optimal stationary base-stock policy, with order-up-to level $S_\alpha^\infty(r)$. In addition, for $L_0 \geq 0$, $L > L_0 + 1$, and $\alpha \in (0, 1)$,

$$OPT(L) \geq c(\mathbb{E}[D] - r_L) + (1 - \alpha) V_\alpha^\infty(r_L, S_\alpha^\infty(r_L)) - U_0 (1 - \alpha)^{-3} L \alpha^{L - L_0}. \quad (9)$$

We now formally define the Markov process representing the inventory position process under such an optimal stationary base-stock policy, initialized in state $S_\alpha^\infty(r_L)$. Let $S_{\alpha, L} \triangleq S_\alpha^\infty(r_L)$. For $r \in [0, \mathbb{E}[D]]$ and $y \in \mathbb{R}$, let $\{X_k^{r, y}, k \geq 1\}$ denote the following Markov process. $X_1^{r, y}$ equals y . For all $k \geq 1$, $X_{k+1}^{r, y} = \max(X_k^{r, y} + r - D_k, y)$. Let $W_k^r \triangleq \sum_{j=1}^k (r - D_j)$, $Z_k^r \triangleq \max_{i \in [0, k-1]} W_i^r$, $Z_\infty^r \triangleq \sup_{i \geq 0} W_i^r$, $M_k^r \triangleq \mathbb{E}[Z_k^r]$, $M_\infty^r \triangleq \mathbb{E}[Z_\infty^r]$. It follows from the well-known analysis of the single-server queue using Lindley's recursion (cf. Asmussen (2003)) that $X_k^{r, y} \sim y + Z_k^r$; and $X_\infty^{r, y} \triangleq \lim_{k \rightarrow \infty} X_k^{r, y}$ (in the sense of weak convergence) is a well-defined r.v. distributed as $y + Z_\infty^r$.

Combining these definitions with Lemmas 3 and 4, we conclude the following.

COROLLARY 3. For $L_0 \geq 0, L > L_0 + 1$, and $\alpha \in (0, 1)$,

$$OPT(L) \geq c(\mathbb{E}[D] - r_L) + (1 - \alpha) \sum_{k=1}^{\infty} \alpha^{k-1} \mathbb{E} \left[G \left(S_{\alpha,L} + Z_k^{r_L} - \sum_{i=1}^{L_0+1} (D'_i - r_L) \right) \right] - U_0(1 - \alpha)^{-3} L \alpha^{L-L_0}.$$

We now briefly review some useful properties of Z_k^r , which we will use to complete the proof of our main results. These properties follow by combining generally well-known results for generating functions, large deviations, single-server queues, and recurrent random walks (cf. Spitzer (1956), Kingman (1962), Folland (1999), Asmussen (2003), Xin and Goldberg (2015)), and we omit the details.

LEMMA 5. For all $r > 0$, $\{M_k^r, k \geq 1\}$ is non-decreasing, $M_\infty^r = \lim_{k \rightarrow \infty} M_k^r$, and for all $i \geq j \geq 1$, $M_i^r - M_j^r = \sum_{k=j}^{i-1} k^{-1} \mathbb{E}[\max(0, W_k^r)]$. If there exists $\epsilon \in (0, \mathbb{E}[D])$ s.t. $r \leq \mathbb{E}[D] - \epsilon$, then $M_\infty^r < \infty$, and $M_\infty^r - M_n^r \leq (\vartheta_\epsilon(1 - \gamma_\epsilon))^{-1} \gamma_\epsilon^n$ for all $n \geq 1$.

Finally, we will also need the following corollary (of Lemma 5), which shows that r_L is uniformly bounded away from $\mathbb{E}[D]$ in an appropriate sense, and whose proof we again defer to the technical appendix Section 5.

COROLLARY 4. For all $L > \epsilon_0^{-2} + L_0 + 1$, it holds that $r_L < \mathbb{E}[D] - \epsilon_0$.

We now complete the proof of our main results.

Proof of Theorem 1 It follows from (3) that for all $\alpha \in (0, 1)$,

$$\begin{aligned} C(\pi_{r_L, S_{\alpha,L} + (L_0+1)r_L}) &= c(\mathbb{E}[D] - r_L) + \mathbb{E} \left[G \left(S_{\alpha,L} + Z_\infty^{r_L} + (L_0 + 1)r_L - \sum_{i=1}^{L_0+1} D'_i \right) \right] \\ &= c(\mathbb{E}[D] - r_L) + (1 - \alpha) \sum_{k=1}^{\infty} \alpha^{k-1} \mathbb{E} \left[G \left(S_{\alpha,L} + Z_\infty^{r_L} - \sum_{i=1}^{L_0+1} (D'_i - r_L) \right) \right]. \end{aligned}$$

Combining with Corollaries 3 and 4, Lemma 5, (1), and the fact that $L > \epsilon_0^{-2} + L_0 + 1$, we conclude that for all $\alpha \in (0, 1)$, $C(\pi_{r^*, S^*}) - OPT(L) - U_0(1 - \alpha)^{-3} L \alpha^{L-L_0}$ is at most

$$\begin{aligned} &(1 - \alpha) \sum_{k=1}^{\infty} \alpha^{k-1} \left(\mathbb{E} \left[G \left(S_{\alpha,L} + Z_\infty^{r_L} - \sum_{i=1}^{L_0+1} (D'_i - r_L) \right) \right] - \mathbb{E} \left[G \left(S_{\alpha,L} + Z_k^{r_L} - \sum_{i=1}^{L_0+1} (D'_i - r_L) \right) \right] \right) \\ &\leq \max(b, h)(1 - \alpha) \sum_{k=1}^{\infty} \alpha^{k-1} (\vartheta_{\epsilon_0}(1 - \gamma_{\epsilon_0}))^{-1} \gamma_{\epsilon_0}^k \\ &= \max(b, h) \gamma_{\epsilon_0} (\vartheta_{\epsilon_0}(1 - \gamma_{\epsilon_0}))^{-1} \frac{1 - \alpha}{1 - \gamma_{\epsilon_0} \alpha} \\ &\leq (1 - \alpha) \max(b, h) \gamma_{\epsilon_0} \vartheta_{\epsilon_0}^{-1} (1 - \gamma_{\epsilon_0})^{-2}. \end{aligned}$$

We conclude that for all $\alpha \in (\frac{1}{2}, 1)$, $C(\pi_{r^*, s^*}) - \text{OPT}(L)$ is at most

$$U_0 2^{L_0} (1 - \alpha)^{-3} L \alpha^L + (1 - \alpha) \max(b, h) \gamma_{\epsilon_0} \vartheta_{\epsilon_0}^{-1} (1 - \gamma_{\epsilon_0})^{-2}.$$

As $L > \epsilon_0^{-2} + L_0 + 1$ implies $L > 100$, which itself may be shown to imply that $5 \frac{\log(L)}{L} < \frac{1}{2}$, we may set $\alpha = 1 - 5 \frac{\log(L)}{L}$. Then applying the fact that $1 - \alpha \leq \exp(-\alpha)$, we conclude that

$$C(\pi_{r^*, s^*}) - \text{OPT}(L) \leq 5 \frac{\log(L)}{L} (U_0 2^{L_0} + \max(b, h) \gamma_{\epsilon_0} \vartheta_{\epsilon_0}^{-1} (1 - \gamma_{\epsilon_0})^{-2}).$$

As $\frac{\log(L)}{L} < L^{-\frac{1}{2}}$ for all $L \geq 1$, combining with the fact that $\text{OPT} \geq g$ and a straightforward calculation completes the proof. \square .

4. Conclusion

In this paper, we proved that when the lead time of the express source is held fixed, a simple TBS policy is asymptotically optimal for the dual-sourcing inventory problem as the lead time of the regular source grows large. Our results provide a solid theoretical foundation for several conjectures and numerical experiments appearing previously in the literature regarding the good empirical performance of such policies. Furthermore, the simple TBS policy performs nearly optimally exactly when standard DP-based methodologies become intractable due to the curse of dimensionality. In addition, since the “best” TBS policy can be computed by solving a convex program that does not depend on the lead time of the regular source, and is easy to implement, our results lead directly to very efficient algorithms with asymptotically optimal performance guarantees. We also explicitly bound the optimality gap of the TBS policy for any fixed lead time (of the regular source), and prove that this decays inverse-polynomially in the lead time of the regular source. Perhaps most importantly, since many companies are already implementing such TBS policies, our results provide strong theoretical support for the widespread use of TBS policies in practice.

This work leaves many interesting directions for future research. First, it would be interesting to further investigate the rate of convergence to optimality of TBS policies as the lead time grows large, especially in light of their use in practical settings. Although we have not optimized the

explicit bounds which we have proven on the optimality gap, we suspect that proving significantly stronger (e.g. exponentially decaying) bounds will require the development of new techniques. For example, when we apply the conditional Jensen’s inequality to lower bound the optimal value by a certain single-sourcing problem in Section 3.1.3, our current approach does not incorporate the fact that $\mathbb{E}[q_i^{*,L}]$ is the same for all i , instead only using the fact that $\mathbb{E}[q_i^{*,L}|D_{[i-1]}]$ is a measurable function of $D_{[i-1]}$. It seems plausible that incorporating this “stationary expectations” property may be a promising approach here. Previous bounds from the literature on the rate of convergence of finite horizon inventory optimization problems to their infinite horizon counterparts, e.g. Hordijk and Tijms (1974, 1975), may also be helpful.

Second, and related to the aforementioned discussion as regards the rate of convergence to optimality of TBS policies, it would be interesting to identify other more sophisticated algorithms which perform better for small-to-moderate lead times, yet remain efficient to implement. Indeed, it remains an interesting open question to better understand the trade-off between algorithmic run-time and achievable performance guarantees in this context, i.e. how complex an algorithm is required to “exploit” the weak correlations which persist even as the lead time grows large. In the context of dual-sourcing, potential algorithms here include: the so-called dual-sourcing smoothing policies recently studied in Boute and Van Mieghem (2015); affine policies more generally (cf. Bertsimas, Iancu and Parrilo (2010)), of which dual-sourcing smoothing policies are a special case; the single index and dual index policies discussed earlier; or the dual-balancing policies analyzed in Levi, Janakiraman and Nagarajan (2008). It would also be quite interesting to analyze “hybrid” algorithms, which could e.g. solve a large dynamic program when the lead time is small, and gradually transition to using simpler heuristics as the lead time grows large; or combine different heuristics depending on the specific problem parameters. In the context of the above conversation on optimality gaps, we do remind the reader that for any fixed regular lead time a TBS policy is not exactly optimal except in some very special cases (cf. JSS), and that our results (and associated insights) should always be applied with care.

On a final note, combined with the results of Goldberg et al. (2015) and Xin and Goldberg (2015), our methodology lays the foundations for a completely new approach to analyzing inventory models with large lead times. So far, this approach has been successful in yielding key insights and efficient algorithms for two settings previously believed intractable: lost sales models with large lead times, and dual-sourcing models with large lead time gap. We believe that our techniques have the potential to make similar progress on many other difficult supply chain optimization problems of practical relevance in which there is a lag between when policy decisions are made and when those decisions are implemented. This includes both more realistic variants of the lost-sales and dual-sourcing models considered so far (e.g. models with distributional dependencies, parameter uncertainty, complex network structure, and more accurate modeling of costs), as well as fundamentally different models (e.g. inventory systems with remanufacturing when the manufactured and remanufactured lead times differ, cf. Zhou, Tao and Chao (2011); multi-echelon systems with lost sales and positive lead times, cf. Huh and Janakiraman (2010); or models with perishable goods). In closing, we note that our approach can more generally be viewed as a methodology to formalize the notion that when there is a high level of uncertainty and randomness in one's supply chain, even simple policies perform nearly as well as very sophisticated policies, since no algorithm can "beat the noise". Exploring this concept from a broader perspective may be fruitful in yielding novel algorithms and insights for a multitude of problems in operations management and operations research.

References

- Allon, G., J. A. Van Mieghem. 2010. Global Dual Sourcing: Tailored Base-Surge Allocation to Near- and Offshore Production. *Management Science*. **56**(1) 110-124.
- Angelus, A., Ö. Özer, 2015. Knowledge you can act on: optimal policies for assembly systems with expediting and advance demand information. working paper.
- Arapostathis, A., V. S. Borkar, E. Fernández-Gaucherand, M. K. Ghosh, S. I. Marcus. 1993. Discrete time controlled Markov processes with average cost criterion: a survey. *SIAM Journal on Control and Optimization*. **31** 282-344.

- Asmussen, S. 2003. *Applied probability and queues*. Second ed., Springer.
- Barankin, E. W. 1961. A delivery-lag inventory model with an emergency provision. *Naval Research Logistics Quarterly*. **8**(3) 285-311.
- Bertsimas, D., D. A. Iancu, P.A. Parrilo. 2010. Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research*. **35**(2) 363-394.
- Billingsley, P. 1999. *Convergence of Probability Measures*. Wiley, New York.
- Boute, R.N., J.A. Van Mieghem. 2015. Global dual sourcing and order smoothing: The impact of capacity and leadtimes. *Management Science*. **61**(9) 2080-2099.
- Chen, Y., W. Xue, J. Yang. 2013. Optimal inventory policy in the presence of a long-term supplier and a spot market. *Operations Research*. **61**(1) 88-97.
- Daniel, K. H. 1963. A delivery-lag inventory model with emergency order. H. Scarf, D. Gilford, M. Shelly, eds. *Multistage Inventory Models and Techniques*, chapter 2. Stanford University Press, Stanford, CA.
- DeCroix, G., J.S. Song, P. Zipkin. 2005. A series system with returns: Stationary analysis. *Operations Research*. **53** 350-362.
- Erdos, P., M. Kac. 1946. On certain limit theorems of the theory of probability. *Bull. Amer. Math. Soc.* **52** 292-302.
- Feinberg, E.A. 2011. Total expected discounted reward MDPs: existence of optimal policies. *Wiley Encyclopedia of Operations Research and Management Science*.
- Feng, Q., S. Sethi, H. Yan, H. Zhang. 2006. Are base-stock policies optimal in inventory problems with multiple delivery modes? *Operations Research*. **54**(4) 801-807.
- Filar, J. 2007. *Controlled Markov chains, graphs and Hamiltonicity*. Now Publishers Inc.
- Fleischmann, M., R. Kuik. 2003. An optimal inventory control with independent stochastic item returns. *European Journal of Operational Research*. **151** 25-37.
- Folland, G. B. 1999. *Real Analysis: Modern Techniques and Their Applications*. Second Edition, Wiley.
- Fox, E.J., R. Metters, J. Semple. 2006. Optimal inventory policy with two suppliers. *Operations Research*. **54**(2) 389-393.

- Fukuda, Y. 1964. Optimal policies for the inventory problem with negotiable leadtime. *Management Science*. **10** 690-708.
- Goldberg, D.A., D. A. Katz-Rogozhnikov, Y. Lu, M. Sharma, M. S. Squillante. 2015. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. To appear in *Mathematics of Operations Research*.
- Gong, X., X. Chao, S. Zheng. 2014. Dynamic pricing and inventory management with dual suppliers of different lead times and disruption risks. *Production and Operations Management*. forthcoming.
- Heyman, D., M. Sobel. 1984. *Stochastic Models in Operations Research*, Vol. II. McGraw-Hill, New York.
- Hordijk, A., H. Tijms. 1974. Convergence Results and Approximations for Optimal (s, S) Policies. *Management Science*. **20**(11) 1432-1438.
- Hordijk, A., H. Tijms. 1975. On a conjecture of Iglehart. *Management Science*. **21**(11) 1342-1345.
- Hua, Z., Y. Yu, W. Zhang, X. Xu. 2014. Structural properties of the optimal policy for dual-sourcing systems with general lead times. *IIE Transactions*. forthcoming.
- Huggins, E. L., T. L. Olsen. 2010. Inventory control with generalized expediting. *Operations Research*. **58**(5) 1414-1426.
- Huh, W.T., G. Janakiraman. 2010. On the optimal policy structure in serial inventory systems with lost sales. *Operations Research*. **58**(2) 486-491.
- Huh, W.T., G. Janakiraman, M. Nagarajan. 2011. Average cost single-stage inventory models: an analysis using a vanishing discount approach. *Operations Research*. **59**(1) 143-155.
- Iglehart, D. L. 1963. Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science*. **9**(2) 259-267.
- Janakiraman, G., S. Seshadri, A. Sheopuri. 2015. Analysis of Tailored Base-Surge Policies in Dual Sourcing Inventory Systems. *Management Science*. **61**(7) 1547-1561.
- Janssen, F., T. De Kok. 1999. A two-supplier inventory model. *International Journal of Production Economics*. **59** 395-403.

- Karlin, S., H. Scarf. 1958. Inventory models of the Arrow-Harris-Marschak type with time lag. K. Arrow, S. Karlin, H. Scarf (Eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford University Press, Stanford, CA, 155-178.
- Kelleher, K. 2003. Why FedEx is gaining ground. *Business 2.0* (October) 56-57.
- Kingman, J.F.C. 1962. Some inequalities for the queue GI/G/1. *Biometrika* **49** 315-324.
- Klosterhalfen, S., G. Kiesmüller, S. Minner. 2011. A comparison of the constant-order and dual-index policy for dual sourcing. *International Journal of Production Economics*. **133**(1) 302-311.
- Korolev, V. Yu, and I. G. Shevtsova. 2010. On the upper bound for the absolute constant in the Berry-Esseen inequality. *Theory of Probability and Its Applications*. **54**(4) 638-658.
- Levi, R., Ganesh Janakiraman, Mahesh Nagarajan. 2008. A 2-approximation algorithm for stochastic inventory control models with lost sales. *Mathematics of Operations Research*. **33**(2) 351-374.
- Minner, S. 2003. Multiple-supplier inventory models in supply chain management: A review. *International Journal of Production Economics*. **81-82** 265-279.
- Morton, T. E. 1969. Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM Review*. **11**(4) 572-596.
- Neuts, F. 1964. An inventory model with optimal time lag. *SIAM Journal on Applied Mathematics*. **12** 179-185.
- Parker, R., R. Kapuscinski. 2004. Optimal policies for a capacitated twoechelon inventory system. *Operations Research*. **52** 739-755.
- Rao, U., A. Scheller-Wolf, S. Tayur. 2000. Development of a rapidresponse supply chain at Caterpillar. *Operations Research*. **48**(2) 189-204.
- Rosenshine, M., D. Obee. 1976. Analysis of a standing order inventory system with emergency orders. *Operations Research*. **24**(6) 1143-1155.
- Rossi, R., W. A. Rijpkema, J. G. A. J. van der Vorst. 2012. The impact of dual sourcing on food supply chain networks: the case of Egyptian strawberries. In *Proc. of the 11th Wageningen International Conference on Chain and Network Management*, Wageningen.

- Scarf, H. 1960. The Optimality of (s, S) Policies in the Dynamic Inventory Problem. *Mathematical Methods in the Social Sciences*. Stanford University Press. 196-202.
- Schäl, M. 1993. Average optimality in dynamic programming with general state space. *Mathematics of Operations Research*. **18**(1) 163-172.
- Scheller-Wolf, A., S. Veeraraghavan, G.-J. van Houtum. 2008. Effective Dual Sourcing with a Single Index Policy. Working paper.
- Sennott, L.I., 1989. Average cost optimal stationary policies in infinite state markov decision processes with unbounded costs. *Operations Research*. **37**(4) 626-633.
- Sheopuri, A., G. Janakiraman, S. Seshadri. 2010. New policies for the stochastic inventory control problem with two supply sources. *Operations Research*. **58**(3) 734-745.
- Song, J., P. Zipkin. 2009. Inventories with Multiple Supply Sources and Networks of Queues with Overflow Bypasses. *Management Science*. **55**(3) 362-372.
- Spitzer, F. 1956. A combinatorial lemma and its application to probability theory. *Transactions of the American Mathematical Society*. **82**(2) 323-339.
- Thorisson, H. 1992. Construction of a stationary regenerative process. *Stochastic processes and their applications*. **42**(2) 237-253.
- Van Mieghem, J. A. 2008. *Operations Strategy: Principles and Practice*. Dynamic Ideas, Belmont, MA.
- Veeraraghavan, S., A. Scheller-Wolf. 2008. Now or later: Dual index policies for capacitated dual sourcing systems. *Operations Research*. **56**(4) 850-864.
- Whittmore, A. S., S. C. Saunders. 1977. Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics*. **32** 293-305.
- Xin, L., D.A. Goldberg. 2015. Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. To appear in *Operations Research*.
- Zhou, S.X., Z. Tao, X. Chao. 2011. Optimal control of inventory systems with multiple types of remanufacturable products. *Manufacturing & Service Operations Management*. **13**(1) 20-34.
- Zipkin, Paul. 2008a. On the structure of lost-sales inventory models. *Operations Research*. **56**(4) 937-944.

Zipkin, P. 2008. Old and new methods for lost-sales inventory systems. *Operations Research*. **56**(5) 1256-1263.

Zipkin, P. 2000. Foundations of inventory management. Vol. 2. New York: McGraw-Hill.

5. Technical Appendix

5.1. Proof of Theorem 2

5.1.1. Overview of proof. Before providing a formal proof, we first provide an intuitive overview, noting that our proof is similar to several proofs in the literature (cf. Xin and Goldberg (2015) and the references therein). We proceed by constructing a sequence of random vectors, one for each sufficiently small $\epsilon > 0$, and later take an appropriate weak limit (which will become the vector satisfying the conditions of the theorem). As in Xin and Goldberg (2015), given $\epsilon > 0$, we will pick a sufficiently large time T_ϵ s.t. the expected performance of an approximately optimal (possibly non-stationary) policy $\pi^{*,\epsilon}$ up to time T_ϵ is “close” to $\text{OPT}(L)$. We then further prove the existence of a time $T_{1,\epsilon}$ “near” T_ϵ s.t. the expedited inventory position and truncated regular pipeline vector (under policy $\pi^{*,\epsilon}$) are “well-behaved” at time $T_{1,\epsilon}$, which will be necessary for our later arguments, as it will allow us to bound the time needed to “clear the system” if one orders nothing from that time onwards. We then construct a “modified policy” and associated Markov chain, which behaves exactly like the expedited inventory position and truncated regular pipeline vector under $\pi^{*,\epsilon}$ on $[1, T_{1,\epsilon}]$, but after that time forces a sequence of ordering decisions which cause the associated inventory position and pipeline vector to re-enter a state distributed as its initial state, at which time the entire process restarts. We note that due to the process being unbounded from below, here the special initialization involving $-\sum_{i=1}^{\hat{G}} D'_{-i}$ will prove useful. This regenerative structure, combined with our careful selection of $T_{1,\epsilon}$, will allow us to apply the theory of regenerative processes to prove the existence of a stationary distribution for the associated Markov chain, which we will prove to satisfy the conditions of an “approximate” version of Theorem 2 (with the approximation error parametrized by ϵ). Taking a weak limit (as $\epsilon \downarrow 0$) of the associated sequence of random vectors yields a random vector satisfying the conditions of Theorem 2, completing the proof.

As all results in this subsection will be stated for a fixed $L_0 \geq 0$ and $L > L_0 + 1$, we assume

these parameters are fixed and suppress any associated notational dependencies. For $\epsilon > 0$, let $\pi^{*,\epsilon}$ denote some (fixed) policy in $\hat{\Pi}$ s.t. $\text{OPT}(L) > C(\pi^{*,\epsilon}) - \frac{\epsilon}{2}$. Let $U_\epsilon \triangleq L^2(\frac{4U}{\epsilon \min(b,h)} + 2\mathbb{E}[D])$, and $U_{2,\epsilon} \triangleq (L + \frac{U_\epsilon}{\mathbb{E}[D]} + 2)(h + b + c)U_\epsilon$. It follows from the definition of limsup that there exists $T_\epsilon > 100(U_{2,\epsilon} + (U + 1)L)\epsilon^{-1}$ s.t. $C(\pi^{*,\epsilon}) > T^{-1} \sum_{t=L_0+1}^T \mathbb{E}[C_t^{\pi^{*,\epsilon}}] - \frac{\epsilon}{2}$ for all $T \geq (1 - \epsilon)T_\epsilon - L$.

5.1.2. Existence of time $T_{1,\epsilon}$, near T_ϵ , at which inventory and pipeline are small in expectation. We first prove that there must exist a time “close to” T_ϵ at which the expedited inventory position and truncated regular pipeline vector (under policy $\pi^{*,\epsilon}$) are “small” (in absolute value) in expectation.

CLAIM 1. For all $\epsilon \in (0, \min(\frac{1}{2}, \frac{U}{2}))$, there exists $T_{1,\epsilon} \in [(1 - \epsilon)T_\epsilon - L, T_\epsilon]$ s.t.

$$C(\pi^{*,\epsilon}) > T_{1,\epsilon}^{-1} \sum_{t=L_0+1}^{T_{1,\epsilon}} \mathbb{E}[C_t^{\pi^{*,\epsilon}}] - \frac{\epsilon}{2}; \quad (10)$$

for all $k \in [0, L - 1]$,

$$\mathbb{E}[\|\hat{I}_{T_{1,\epsilon}+k-L_0}^{\pi^{*,\epsilon}} + q_{T_{1,\epsilon}+k-L_0}^{\pi^{*,\epsilon},E}\|] \leq \frac{2UL}{\epsilon \min(b,h)} + (L_0 + 1)\mathbb{E}[D]; \quad (11)$$

and for all $k \in [1, L - L_0 - 1]$,

$$\mathbb{E}[\mathcal{R}_k^{\pi^{*,\epsilon}, T_{1,\epsilon}-L_0}] \leq \frac{4UL}{\epsilon \min(b,h)} + 2L\mathbb{E}[D]. \quad (12)$$

Proof of Claim 1 Note that we may (deterministically) partition the time interval $[(1 - \epsilon)T_\epsilon - L, T_\epsilon]$ into $\lceil \frac{\epsilon T_\epsilon}{L} \rceil$ disjoint intervals each of length L , plus an additional disjoint time interval of length possibly less than L . Suppose for contradiction that of these $\lceil \frac{\epsilon T_\epsilon}{L} \rceil$ disjoint time intervals of length L , there does not exist a single such interval I s.t.

$$\mathbb{E}[\|\hat{I}_{t-L_0}^{\pi^{*,\epsilon}} + q_{t-L_0}^{\pi^{*,\epsilon},E}\|] \leq \frac{2UL}{\epsilon \min(b,h)} + (L_0 + 1)\mathbb{E}[D] \text{ for all } t \in I. \quad (13)$$

In that case, by the triangle inequality, each of these $\lceil \frac{\epsilon T_\epsilon}{L} \rceil$ intervals contains at least one time period t for which

$$\mathbb{E}[\|\hat{I}_{t-L_0}^{\pi^{*,\epsilon}} + q_{t-L_0}^{\pi^{*,\epsilon},E} - \sum_{i=t-L_0}^t D_i\|] > \frac{2UL}{\epsilon \min(b,h)}.$$

Hence by (1), non-negativity of costs, and the definition of T_ϵ we conclude that

$$\begin{aligned} C(\pi^{*,\epsilon}) &> \frac{\sum_{t=\lceil(1-\epsilon)T_\epsilon\rceil-L}^{T_\epsilon} \mathbb{E}[C_t^{\pi^{*,\epsilon}}]}{T_\epsilon} - \frac{\epsilon}{2} \\ &> \frac{\min(b, h) \times \frac{\epsilon T_\epsilon}{L} \times \frac{2UL}{\epsilon \min(b, h)}}{T_\epsilon} - \frac{\epsilon}{2} = 2U - \frac{\epsilon}{2} > \frac{3}{2}U, \end{aligned}$$

and thus $\text{OPT}(L) > \frac{3}{2}U - \frac{\epsilon}{2} > U$, a contradiction. Let t' denote the left end-point of the corresponding interval satisfying (13), whose existence we have just proven by contradiction (in case of multiple such intervals, take the left-most such interval). Now, further suppose for contradiction that there exists $k \in [1, L - L_0 - 1]$ s.t. $\mathbb{E}[\mathcal{R}_k^{\pi^{*,\epsilon}, t' - L_0}] > \frac{4UL}{\epsilon \min(b, h)} + 2L\mathbb{E}[D]$. Then it would follow from the inventory update dynamics, non-negativity of order quantities, and the triangle inequality that $\mathbb{E}[\hat{I}_{t'+k-L_0}^{\pi^{*,\epsilon}} + q_{t'+k-L_0}^{\pi^{*,\epsilon}, E}] > \frac{2UL}{\epsilon \min(b, h)} + (L_0 + 1)\mathbb{E}[D]$, which would itself contradict the definition of t' . Combining the above, and setting $T_{1,\epsilon} = t'$, completes the proof. \square .

5.1.3. Statement of approximate form of Theorem 2. We now formally state the aforementioned approximate version of Theorem 2.

LEMMA 6. For all $\epsilon \in (0, \min(\frac{1}{2}, U))$, one may construct an $L - L_0 - 1$ -dimensional random vector $\chi^{*,\epsilon}$, an $L - L_0$ -dimensional random vector $\mathbf{q}^{*,\epsilon}$, and a random variable $\mathcal{I}^{*,\epsilon}$, as well as $\{D_i, i \geq 1\}$, on a common probability space s.t. the following are true.

- (i) $(\chi^{*,\epsilon}, \mathbf{q}^{*,\epsilon}, \mathcal{I}^{*,\epsilon})$ has finite mean, and w.p.1 $(\chi^{*,\epsilon}, \mathbf{q}^{*,\epsilon})$ is non-negative. Also, $(\chi^{*,\epsilon}, \mathcal{I}^{*,\epsilon})$ is independent of $\{D_i, i \geq 1\}$ and $q_i^{*,\epsilon}$ is independent of $\{D_j, j \geq i\}$ for $i \in [1, L - L_0]$.
- (ii) $\chi_i^{*,\epsilon} \sim \chi_1^{*,\epsilon}$ for $i \in [1, L - L_0 - 1]$, and $q_i^{*,\epsilon} \sim q_1^{*,\epsilon}$ for $i \in [1, L - L_0]$.
- (iii) $\mathbb{E}[\chi_1^{*,\epsilon}] + \mathbb{E}[q_1^{*,\epsilon}] = \mathbb{E}[D]$.
- (iv) For all $k \in [1, L - L_0]$,

$$\mathcal{I}^{*,\epsilon} + \sum_{i=1}^{k-1} (q_i^{*,\epsilon} + \chi_i^{*,\epsilon} - D_i) + q_k^{*,\epsilon} - \sum_{i=k}^{k+L_0} D_i \sim \mathcal{I}^{*,\epsilon} + q_1^{*,\epsilon} - \sum_{i=1}^{L_0+1} D_i.$$

(v)

$$\text{OPT}(L) > c\mathbb{E}[q_1^{*,\epsilon}] + \mathbb{E}\left[G\left(\mathcal{I}^{*,\epsilon} + q_1^{*,\epsilon} - \sum_{i=1}^{L_0+1} D_i\right)\right] - \epsilon.$$

5.1.4. Proof of Lemma 6 by construction of a Markov chain with an appropriate stationary distribution. We now construct an appropriate Markov chain which repeatedly mimics $\pi^{*,\epsilon}$ for blocks of time of length $T_{1,\epsilon}$, and then (by a sequence of ordering decisions) brings the system back to a state distributed as its initial state (involving $-\sum_{i=1}^{\hat{G}} D'_{-i}$). This is accomplished by allowing for an extra “time-accounting” dimension in the state-space. While this “clock” is between 1 and $T_{1,\epsilon}$, the Markov chain dynamics parallel those of the inventory and pipeline in $\pi^{*,\epsilon}$ on $[1, T_{1,\epsilon}]$, and the clock increases by one in each period. Whenever the clock reaches $T_{1,\epsilon}$, the Markov chain dynamics instead parallel those of a policy which first orders nothing until the truncated regular pipeline vector clears and the inventory position goes below 0, then places an expedited order to bring the inventory position to exactly 0, and finally orders nothing for an additional geometrically distributed number of time periods, where this geometric idling will preclude any pathological periodicity that might otherwise arise (ensuring existence of relevant stationary measures). This brings the system back to a state in which the truncated regular pipeline vector is empty and the inventory position is distributed as $-\sum_{i=1}^{\hat{G}} D'_{-i}$, at which time the clock restarts to 1 and the cycle repeats, which thus yields a regenerative process. We further note that in the associated Markov chain we will also keep track of the most recent expedited order, so that all relevant inventory and ordering costs can be expressed directly as a function of the state in the associated Markov chain. This will allow us to apply the theory of regenerative processes to prove that the expected value of an appropriate function of the corresponding steady-state vector bounds the average cost incurred by $\pi^{*,\epsilon}$ on $[1, T_{1,\epsilon}]$, which itself well-approximates $\text{OPT}(L)$ (i.e. ensuring that Lemma 6.(v) is satisfied). Combining the above will allow us to prove that this steady-state vector satisfies the conditions of Lemma 6.

Proof of Lemma 6 We construct an $(L - L_0 + 2)$ -dimensional discrete-time Markov process $\{\mathbf{Y}_t^\epsilon, t \geq 1\} = \{(\chi^{\epsilon,t}, q_t^\epsilon, \mathcal{I}_t^\epsilon, \tau_t^\epsilon), t \geq 1\}$, where $\chi^{\epsilon,t}$ is an $(L - L_0 - 1)$ -dimensional random vector, and $q_t^\epsilon, \mathcal{I}_t^\epsilon$, and τ_t^ϵ are random variables. Let $\{B_t, t \geq 1\}$ denote an i.i.d. sequence of Bernoulli r.v.s, each of which equals 1 w.p. $\frac{1}{2}$ and 0 w.p. $\frac{1}{2}$. Then $\{\mathbf{Y}_t^\epsilon, t \geq 1\}$ evolves as follows.

- $\chi^{\epsilon,1} = \mathbf{0}$, $\mathcal{I}_1^\epsilon = -\sum_{i=1}^{\hat{G}} D'_{-i}$, $q_1^\epsilon = \hat{f}_{E,1}^{\pi^{*,\epsilon}}(\chi^{\epsilon,1}, \mathcal{I}_1^\epsilon)$, $\tau_1^\epsilon = 1$.

For $t \geq 1$, the dynamics are as follows.

- $\chi_i^{\epsilon,t+1} = \chi_{i+1}^{\epsilon,t}$ for $i \in [1, L - L_0 - 2]$, $\mathcal{I}_{t+1}^\epsilon = \mathcal{I}_t^\epsilon + \chi_1^{\epsilon,t} + q_t^\epsilon - D_t$.
- If $\tau_t^\epsilon \in [1, T_{1,\epsilon})$:
 - $\tau_{t+1}^\epsilon = \tau_t^\epsilon + 1$, $\chi_{L-L_0-1}^{\epsilon,t+1} = \hat{f}_{R,\tau_t^\epsilon}^{\pi^{*,\epsilon}}(\chi^{\epsilon,t}, \mathcal{I}_t^\epsilon)$, $q_{t+1}^\epsilon = \hat{f}_{E,\tau_{t+1}^\epsilon}^{\pi^{*,\epsilon}}(\chi^{\epsilon,t+1}, \mathcal{I}_{t+1}^\epsilon)$.
- If $\tau_t^\epsilon = T_{1,\epsilon}$ and either $\chi^{\epsilon,t} \neq \mathbf{0}$ or $\mathcal{I}_1^\epsilon > 0$:
 - $\chi_{L-L_0-1}^{\epsilon,t+1} = q_{t+1}^\epsilon = 0$, $\tau_{t+1}^\epsilon = T_{1,\epsilon}$.
- If $\tau_t^\epsilon = T_{1,\epsilon}$ and $\chi^{\epsilon,t} = \mathbf{0}$ and $\mathcal{I}_{t+1}^\epsilon \leq 0$:
 - $\chi_{L-L_0-1}^{\epsilon,t+1} = 0$, $q_{t+1}^\epsilon = -\mathcal{I}_{t+1}^\epsilon$, $\tau_{t+1}^\epsilon = 0$.
- If $\tau_t^\epsilon = 0$:
 - $\chi_{L-L_0-1}^{\epsilon,t+1} = 0$, $q_{t+1}^\epsilon = 0$, $\tau_{t+1}^\epsilon = B_t$.

One may easily verify the following properties of $\{\mathbf{Y}_t^\epsilon, t \geq 1\}$. Let $z(x, y) \triangleq \mathbb{E}[G(x + y - \sum_{i=1}^{L_0+1} D'_i)] + cy$, and \hat{T}_ϵ denote a r.v. distributed as the time between the chain's initial and second visit to a state s.t. $\tau_t^\epsilon = 1$.

- It follows directly from the Markov chain dynamics that for all $t \geq 1$, $\chi_i^{\epsilon,t+1} \sim \chi_{i+1}^{\epsilon,t}$ for $i \in [1, L - 1]$.
- Conditional on the event $\{\tau_t^\epsilon = T_{1,\epsilon}\}$, the expected number of time steps until $\tau_t^\epsilon = 0$ is at most $L + \frac{U_\epsilon}{\mathbb{E}[D]}$.
- Conditional on the event $\{\tau_t^\epsilon = 0, \tau_{t-1}^\epsilon = T_{1,\epsilon}\}$, it holds that (w.p.1) $\chi^{\epsilon,t+1} = \mathbf{0}$, $\mathcal{I}_{t+1}^\epsilon = -D_t$, and the number of time steps until $\tau_t^\epsilon = 1$ is distributed as \hat{G} .
- Conditional on the event $\{\tau_t^\epsilon = 1\}$, it holds that $\mathbf{Y}_t^\epsilon \sim (\mathbf{0}, 0, -\sum_{i=1}^{\hat{G}} D'_{-i}, 1)$, and the joint distribution of $\{\mathbf{Y}_i^\epsilon, i \in [t, t + T_\epsilon - 1]\}$ is identical to that of $\{(\mathcal{R}^{\pi^{*,\epsilon},i}, q_i^{\pi^{*,\epsilon},E}, \hat{I}_i^{\pi^{*,\epsilon}}, i), i \in [1, T_{1,\epsilon}]\}$.
- W.p.1 $\hat{T}_\epsilon \geq T_{1,\epsilon}$, and $\mathbb{E}[\hat{T}_\epsilon] - T_{1,\epsilon} \leq L + \frac{U_\epsilon}{\mathbb{E}[D]} + 2$.
- $0 \leq \mathbb{E}[\sum_{t=1}^{\hat{T}_\epsilon} z(\mathcal{I}_t^\epsilon, q_t^\epsilon)] - \mathbb{E}[\sum_{t=1}^{T_{1,\epsilon}} z(\mathcal{I}_t^\epsilon, q_t^\epsilon)] \leq U_{2,\epsilon}$.

Combining with the basic definitions associated with the theory of regenerative processes (here we refer the interested reader to Asmussen (2003) for an excellent overview), we conclude that

$\{\mathbf{Y}_t^\epsilon, t \geq 1\}$ is a discrete-time aperiodic regenerative process, with regeneration points coinciding with visits to states s.t. $\tau_t^\epsilon = 1$. Then we may conclude the following from standard results in the theory of regenerative processes (cf. Asmussen (2003), Thorisson (1992)).

- (a) $\{\mathbf{Y}_t^\epsilon, t \geq 1\}$ converges weakly (as $t \rightarrow \infty$) to a limiting random vector $\mathbf{Y}_\infty^\epsilon = (\chi^{\epsilon, \infty}, q_\infty^\epsilon, \mathcal{I}_\infty^\epsilon, \tau_\infty^\epsilon)$.
- (b) Initializing the relevant Markov chain with initial conditions distributed as $\mathbf{Y}_\infty^\epsilon$ yields a stationary Markov process $\{\bar{\mathbf{Y}}_t^\epsilon, t \geq 1\} = \{(\bar{\chi}^{\epsilon, t}, \bar{q}_t^\epsilon, \bar{\mathcal{I}}_t^\epsilon, \bar{\tau}_t^\epsilon, t \geq 1\}$. Furthermore, it follows directly from the relevant Markov chain dynamics that we may construct $\{\bar{\mathbf{Y}}_t^\epsilon, t \geq 1\}$ and $\{D_i, i \geq 1\}$ on an appropriate probability space s.t. setting $\chi^{*, \epsilon} = \bar{\chi}^{\epsilon, 1}$, $\mathcal{I}^{*, \epsilon} = \bar{\mathcal{I}}_1^\epsilon$, and $q_k^{*, \epsilon} = \bar{q}_k^\epsilon$ for $k \in [1, L - L_0]$ yields a random vector satisfying conditions (i) - (iv) of Lemma 6.

$$(c) \mathbb{E}[z(\mathcal{I}_\infty^\epsilon, q_\infty^\epsilon)] = \frac{\mathbb{E}[\sum_{t=1}^{\hat{T}_\epsilon} z(\mathcal{I}_t^\epsilon, q_t^\epsilon)]}{\mathbb{E}[\hat{T}_\epsilon]}.$$

Further combining (c) with our previous bounds for $\mathbb{E}[\hat{T}_\epsilon]$, $\mathbb{E}[\sum_{t=1}^{\hat{T}_\epsilon} z(\mathcal{I}_t^\epsilon, q_t^\epsilon)]$, our definition of $T_{1, \epsilon}$, and some straightforward algebra (the details of which we omit) demonstrates that the same random vector exhibited in (b) also satisfies condition (v) of Lemma 6, completing the proof of the lemma. \square .

5.1.5. Proof of Theorem 2. We now complete the proof of Theorem 2, by taking an appropriate weak limit (as $\epsilon \downarrow 0$) of the random vectors which we have proven to satisfy the conditions of Lemma 6, and verifying certain interchanges of expectation and limit (in inequality form).

Proof of Theorem 2 To complete the proof of Theorem 2, we now prove that the sequence of random vectors $\{(\chi^{*, \frac{1}{n}}, \mathbf{q}^{*, \frac{1}{n}}, \mathcal{I}^{*, \frac{1}{n}}), n \geq 2 + \frac{1}{U}\}$ is tight. It follows from Lemma 6.(iii) and (v), non-negativity, the triangle inequality, the fact that $\text{OPT}(L) \leq U$, and (1) that for all $n \geq 2 + \frac{1}{U}$,

$$\mathbb{E}[|\mathcal{I}^{*, \frac{1}{n}}|] \leq \frac{2U}{\min(b, h)} + (L_0 + 1)\mathbb{E}[D]. \quad (14)$$

Combining with Lemma 6.(ii) - (iii) and non-negativity, we conclude the desired tightness, and hence existence of at least one subsequential limit (cf. Billingsley (1999)) $(\chi^{*, \infty}, \mathbf{q}^{*, \infty}, \mathcal{I}^{*, \infty})$. Let $\{n_i, i \geq 1\}$ denote any fixed subsequence along which the sequence of measures converges to this limit, s.t. $n_1 \geq 2 + \frac{1}{U}$. That this weak limit satisfies Theorem 2.(i) - (iii) follows from the definition of weak convergence. However, it will require somewhat subtle reasoning to prove (iv) - (vi),

since e.g. dominated convergence does not necessarily hold and thus one must take care when interchanging limit and expectation. Note that by the Skorohod representation theorem and continuous mapping theorem, we may construct $\{(\chi^{*,\frac{1}{n_i}}, \mathbf{q}^{*,\frac{1}{n_i}}, |\mathcal{I}^{*,\frac{1}{n_i}}|), i \geq 1\}$ and $(\chi^{*,\infty}, \mathbf{q}^{*,\infty}, |\mathcal{I}^{*,\infty}|)$ on a common probability space s.t. the corresponding sequence of random vectors converges almost surely (i.e. not only in distribution) to $(\chi^{*,\infty}, \mathbf{q}^{*,\infty}, |\mathcal{I}^{*,\infty}|)$. As all associated r.v.s are non-negative, we may apply Fatou's lemma to conclude that $\mathbb{E}[\chi_1^{*,\infty}] \leq \liminf_{i \rightarrow \infty} \mathbb{E}[\chi_1^{*,\frac{1}{n_i}}]$, $\mathbb{E}[q_1^{*,\infty}] \leq \liminf_{i \rightarrow \infty} \mathbb{E}[q_1^{*,\frac{1}{n_i}}]$, and $\mathbb{E}[|\mathcal{I}^{*,\infty}|] \leq \liminf_{i \rightarrow \infty} \mathbb{E}[|\mathcal{I}^{*,\frac{1}{n_i}}|]$. Combining with Lemma 6.(i) and (iii), as well as (14), then completes the proof of Theorem 2.(iv). Combining with the already proven Theorem 2.(ii) and (iii), with $k=2$, yields Theorem 2.(v). Finally, we prove that the corresponding vector also satisfies Theorem 2.(vi). Let $Z_n \triangleq cq_1^{*,\frac{1}{n}} + G\left(\mathcal{I}^{*,\frac{1}{n}} + q_1^{*,\frac{1}{n}} - \sum_{i=1}^{L_0+1} D_i\right)$, and $Z_\infty \triangleq cq_1^{*,\infty} + G\left(\mathcal{I}^{*,\infty} + q_1^{*,\infty} - \sum_{i=1}^{L_0+1} D_i\right)$. The already proven weak convergence, and continuous mapping theorem, implies that $\{Z_{n_i}, i \geq 1\}$ converges weakly to Z_∞ . It follows from the Skorohod representation theorem (cf. Billingsley (1999)) that we may construct $\{Z_{n_i}, i \geq 1\}$ and Z_∞ on a common probability space so that this convergence holds almost surely (as opposed to only in distribution). Applying non-negativity and Fatou's lemma, we conclude that on this probability space, $\mathbb{E}[\liminf_{i \rightarrow \infty} Z_{n_i}] \leq \liminf_{i \rightarrow \infty} \mathbb{E}[Z_{n_i}]$, and hence (combining with the stated almost sure convergence) $\mathbb{E}[Z_\infty] \leq \liminf_{i \rightarrow \infty} \mathbb{E}[Z_{n_i}]$. Combining with Lemma 6.(v), which implies that $\text{OPT}(L) > \mathbb{E}[Z_{n_i}] - \frac{1}{n_i}$ for all $i \geq 1$, and the already proven Theorem 2.(v), completes the proof. \square .

5.2. Proof of Lemma 4

In preparation for bounding $V_\alpha^\infty(r, x) - V_\alpha^n(r, x)$, we first bound the optimal value, and set of minimizers, of $V_\alpha^n(r, x)$, uniformly in n . For $\alpha \in (0, 1)$ and $r \in \mathbb{R}$, let $\overline{S}_\alpha^n(r)$ denote the supremum of the set of minimizers (with respect to x) of $V_\alpha^n(r, x)$, where we note that a straightforward contradiction demonstrates that $\overline{S}_\alpha^n(r) \in (-\infty, \infty)$ for each α, n, r ; and it follows from Lemma 3 that $V_\alpha^n(r, -\infty) = V_\alpha^n(r, \overline{S}_\alpha^n(r))$. Then we prove the following uniform bounds.

LEMMA 7. 1. For $\alpha \in (0, 1)$ and $r, x \in \mathbb{R}$, it holds that

$$\sup_{n \geq 1} V_\alpha^n(r, x) < 2(L_0 + 1) \max(b, h)(|x| + |r| + \mathbb{E}[D])(1 - \alpha)^{-2}.$$

2. For $\alpha \in (0, 1)$ and $n \geq 1$, it holds that $|\bar{S}_\alpha^n(r)| < \bar{S}_\alpha(r)$.

3. For all $y \notin [-\bar{S}_\alpha(r), \bar{S}_\alpha(r)]$ and $n \geq 1$,

$$\mathbb{E} \left[G(y - \sum_{k=n}^{L_0+n} (D_k - r)) \right] + \alpha \mathbb{E} [V_\alpha^{n-1}(r, y - (D_{L_0+n} - r))] > V_\alpha^n(r, \bar{S}_\alpha^n(r)) + (L_0 + 1) \max(b, h) \mathbb{E}[D].$$

4. For all $L > L_0 + 1$,

$$OPT(L) \geq c(\mathbb{E}[D] - r_L) + (1 - \alpha)V_\alpha^{L-L_0}(r_L, -\bar{S}_\alpha(r_L)).$$

Proof of Lemma 7 By evaluating the policy which never orders, we conclude that for all $\alpha \in (0, 1)$, $r, x \in \mathbb{R}$, $\sup_{n \geq 1} V_\alpha^n(r, x)$ is at most $\mathbb{E} \left[\sum_{i=1}^{\infty} \alpha^{i-1} G(x - \sum_{j=1}^i (D_j - r) - \sum_{k=i+1}^{L_0+i} (D_k - r)) \right]$, which by (1) is itself bounded by

$$\max(b, h)(|x| + |r| + \mathbb{E}[D]) \sum_{i=1}^{\infty} (i + L_0) \alpha^{i-1} < 2(L_0 + 1) \max(b, h)(|x| + |r| + \mathbb{E}[D])(1 - \alpha)^{-2}.$$

The remainder of the lemma follows from (1), Lemmas 2 and 3, and a straightforward calculation and argument by contradiction, and we omit the details. Combining the above completes the proof.

□

With Lemma 7 in hand, we now complete the proof of Lemma 4.

Proof of Lemma 4 We first demonstrate that $V_\alpha^\infty(r, x) = \lim_{n \rightarrow \infty} V_\alpha^n(r, x)$, and complete the proof of (8). The existence of the corresponding limit follows from the monotonicity (in n) guaranteed by Lemma 3. That $V_\alpha^\infty(r, x) \geq \lim_{n \rightarrow \infty} V_\alpha^n(r, x)$ for all $\alpha \in (0, 1)$ and $r, x \in \mathbb{R}$ follows immediately from the definitions of the associated optimization problems. To prove the other direction, as well as (8), we note that for any fixed $n \geq 1$, it follows from the convexity ensured by Lemma 3 that there exists an optimal policy π for the problem stated in the r.h.s. of (5) of base-stock form, with order-up-to levels C_1, \dots, C_n (i.e. order up to level C_i in period i if the pre-order inventory level is below C_i , otherwise order nothing). Furthermore, it follows from Lemma 7 that $\max_{i=1, \dots, n} |C_i| \leq \bar{S}_\alpha(r)$. Now, consider the policy π' (for the problem stated in the r.h.s. of (6)) that orders up to level C_i in period i if the pre-order inventory position is below C_i and otherwise orders nothing in periods $i = 1, \dots, n$; and orders nothing in all remaining periods, irregardless of the inventory level.

It follows from a straightforward bounding argument that under policy π' , w.p.1 the absolute value of the post-ordering inventory position in period i is at most $|x| + \bar{S}_\alpha(r) + (i-1)|r| + \sum_{k=1}^{i-1} D_k$. Thus by the dynamics of the underlying inventory problem and (1), it follows that for all $i \geq n+1$, $C_i^{\pi'} \leq \max(b, h)(|x| + \bar{S}_\alpha(r) + (i+L_0)|r| + (i+L_0)\mathbb{E}[D])$. Thus since (by construction) $C_i^{\pi'} = C_i^\pi$ for $i \in [1, n]$, it follows from definitions and straightforward algebra that

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^{\infty} \alpha^{i-1} C_i^{\pi'}(r, x) \right] - V_\alpha^n(r, x) &\leq \max(b, h)(\bar{S}_\alpha(r) + |x| + |r| + \mathbb{E}[D]) \sum_{i=n+1}^{\infty} (i+L_0) \alpha^{i-1} \\ &\leq \max(b, h)(\bar{S}_\alpha(r) + |x| + |r| + \mathbb{E}[D])(1+L_0+n)(1-\alpha)^{-2} \alpha^n. \end{aligned}$$

This completes the proof of (8), and letting $n \rightarrow \infty$ completes the proof that $V_\alpha^\infty(r, x) = \lim_{n \rightarrow \infty} V_\alpha^n(r, x)$. Combining with Lemmas 3 and 7, the fact that convexity and monotonicity are preserved under limits, and a straightforward contradiction argument completes the proof of all parts of the lemma regarding properties of $V_\alpha^\infty(r, x)$ and the associated optimization problems and optimal policies. Finally, we complete the proof of (9). It follows from Lemmas 3 and 7, the already proven parts of Lemma 4, and the fact that Theorem 2 ensures $r_L \in [0, \mathbb{E}[D]]$ that $\text{OPT}(L) - c(\mathbb{E}[D] - r_L)$ is at least

$$(1-\alpha) \left(V_\alpha^\infty(r_L, S_\alpha^\infty(r_L)) - 2 \max(b, h)(\bar{S}_\alpha(\mathbb{E}[D]) + \mathbb{E}[D])(1+L)(1-\alpha)^{-2} \alpha^{L-L_0} \right).$$

Combining with some straightforward algebra, the definition of U_0 , and the already proven parts of Lemma 4 completes the proof. \square .

5.3. Proof of Corollary 4

Before proving Corollary 4, we will need a preliminary result which demonstrates that if r is “very close” to $\mathbb{E}[D]$, then M_i^r is “very large” for an appropriate range of i . We will then use this result to show that r_L cannot be “too close” to $\mathbb{E}[D]$ by deriving a contradiction, showing that in this case the optimal value would be strictly greater than U , which is impossible.

LEMMA 8. *If there exists $\epsilon \in [0, \mathbb{E}[D] - Q_0]$ s.t. $r \in (\mathbb{E}[D] - \epsilon, \mathbb{E}[D])$, then for all $i, j \in [400\hat{p}_0^{-2}, (\hat{p}_0\eta_0\epsilon^{-1})^2]$ s.t. $i \geq j$,*

$$M_i^r - M_j^r \geq \frac{1}{5} \hat{p}_0 \eta_0 (i^{\frac{1}{2}} - j^{\frac{1}{2}}) - (i-j)\epsilon - 2\eta_0 \left(\log\left(\frac{i}{j}\right) + 2 \right).$$

Proof of Lemma 8 We note that the result would follow from well-known weak-convergence results under additional assumptions on D (e.g. finite variance, cf. Erdos and Kac (1946)), but to avoid unnecessary assumptions (and for completeness) we provide a proof from first principles. Let us fix any $\epsilon \in [0, \mathbb{E}[D] - Q_0]$, $r \in (\mathbb{E}[D] - \epsilon, \mathbb{E}[D]]$, and $k \in [400\hat{p}_0^{-2}, (\hat{p}_0\eta_0\epsilon^{-1})^2]$ (supposing this interval is non-empty, i.e. $\epsilon \leq \frac{\hat{p}_0^2\eta_0}{20}$). Let $\{A_i^{+,r}, i \geq 1\}$ denote an i.i.d. sequence of r.v.s distributed as $r - D$ conditioned on the event $\{r > D\}$, and $\{A_i^{-,r}, i \geq 1\}$ denote an i.i.d. sequence of r.v.s distributed as $D - r$ conditioned on the event $\{D \geq r\}$. Let B_k^r denote a binomially distributed r.v. with parameters $k, \rho_r \triangleq \mathbb{P}(\{r > D\})$, independent of $\{A_i^{+,r}, i \geq 1\}$ and $\{A_i^{-,r}, i \geq 1\}$. It follows from definitions and the constraints on ϵ and r that

$$\rho_r \in [\frac{1}{2}p_0, p_0].$$

Note that for $k \geq 1$, we may construct W_k^r on an appropriate probability space s.t. $W_k^r = \sum_{i=1}^{B_k^r} A_i^{+,r} - \sum_{i=1}^{k-B_k^r} A_i^{-,r}$, in which case (by non-negativity) $\mathbb{E}[\max(0, W_k^r)]$ is at least

$$\mathbb{E} \left[\sum_{i=1}^{B_k^r} A_i^{+,r} - \sum_{i=1}^{k-B_k^r} A_i^{-,r} \middle| \left\{ B_k^r \geq \rho_r k + (\rho_r(1-\rho_r)k)^{\frac{1}{2}} \right\} \right] \mathbb{P} \left(\frac{B_k^r - \rho_r k}{(\rho_r(1-\rho_r)k)^{\frac{1}{2}}} \geq 1 \right).$$

Furthermore, since $\rho_r \mathbb{E}[A_1^{+,r}] = (1 - \rho_r) \mathbb{E}[A_1^{-,r}] - (\mathbb{E}[D] - r)$, it follows from non-negativity and independence that

$$\begin{aligned} & \mathbb{E} \left[\sum_{i=1}^{B_k^r} A_i^{+,r} - \sum_{i=1}^{k-B_k^r} A_i^{-,r} \middle| \left\{ B_k^r \geq \rho_r k + (\rho_r(1-\rho_r)k)^{\frac{1}{2}} \right\} \right] \\ & \geq (\rho_r k + (\rho_r(1-\rho_r)k)^{\frac{1}{2}}) \mathbb{E}[A_1^{+,r}] - ((1-\rho_r)k - (\rho_r(1-\rho_r)k)^{\frac{1}{2}}) \mathbb{E}[A_1^{-,r}] \\ & = (\rho_r(1-\rho_r)k)^{\frac{1}{2}} (\mathbb{E}[A_1^{+,r}] + \mathbb{E}[A_1^{-,r}]) - k(\mathbb{E}[D] - r). \end{aligned}$$

Let $N(0,1)$ denote a standard normal r.v. By the celebrated Berry-Esseen Theorem (cf. Korolev et al. (2010)),

$$\left| \mathbb{P} \left(\frac{B_k^r - \rho_r k}{(\rho_r(1-\rho_r)k)^{\frac{1}{2}}} \geq 1 \right) - \mathbb{P}(N(0,1) \geq 1) \right| \leq 2(\rho_r(1-\rho_r)k)^{-\frac{1}{2}}.$$

It is easily verified from definitions that

$$\mathbb{E}[A_1^{+,r}] + \mathbb{E}[A_1^{-,r}] \geq \eta_0 \quad , \quad (\rho_r(1-\rho_r))^{\frac{1}{2}} \geq \hat{p}_0 \quad , \quad \mathbb{P}(N(0,1) \geq 1) \geq \frac{1}{10}.$$

Thus combining the above, we conclude that

$$\mathbb{E} \left[\sum_{i=1}^{B_k^r} A_i^{+,r} - \sum_{i=1}^{k-B_k^r} A_i^{-,r} \middle| \left\{ B_k^r \geq \rho_r k + (\rho_r(1-\rho_r)k)^{\frac{1}{2}} \right\} \right] \geq \hat{p}_0 \eta_0 k^{\frac{1}{2}} - \epsilon k, \quad (15)$$

and

$$\mathbb{P} \left(\frac{B_k^r - \rho_r k}{(\rho_r(1-\rho_r)k)^{\frac{1}{2}}} \geq 1 \right) \geq \frac{1}{10} - 2\hat{p}_0^{-1} k^{-\frac{1}{2}}. \quad (16)$$

As our assumptions on ϵ, r, k ensure that the r.h.s. of both (15) and (16) are non-negative, we conclude that $\mathbb{E}[\max(0, W_k^r)]$ is at least $(\hat{p}_0 \eta_0 k^{\frac{1}{2}} - \epsilon k) \left(\frac{1}{10} - 2\hat{p}_0^{-1} k^{-\frac{1}{2}} \right)$, which is itself at least $\frac{1}{10} \hat{p}_0 \eta_0 k^{\frac{1}{2}} - k\epsilon - 2\eta_0$. Thus by Lemma 5,

$$\begin{aligned} M_i - M_j &= \sum_{l=j}^{i-1} l^{-1} \mathbb{E}[\max(0, W_l^r)] \\ &\geq \frac{1}{10} \hat{p}_0 \eta_0 \sum_{l=j}^{i-1} l^{-\frac{1}{2}} - (i-j)\epsilon - 2\eta_0 \sum_{l=j}^{i-1} l^{-1} \\ &\geq \frac{1}{10} \hat{p}_0 \eta_0 \int_j^i x^{-\frac{1}{2}} dx - (i-j)\epsilon - 2\eta_0 \left(\log\left(\frac{i}{j}\right) + 2 \right) \\ &= \frac{1}{5} \hat{p}_0 \eta_0 \left(i^{\frac{1}{2}} - j^{\frac{1}{2}} \right) - (i-j)\epsilon - 2\eta_0 \left(\log\left(\frac{i}{j}\right) + 2 \right), \end{aligned}$$

where we have used the well-known fact that for all $n \geq 1$, $\log(n) \leq \sum_{l=1}^n l^{-1} \leq \log(n) + 2$. Combining the above completes the proof. \square

Before completing the proof of Corollary 4, it will be useful to collect a few additional auxiliary bounds. For $\alpha \in (0, 1)$, let G_α denote a geometrically distributed r.v. with success probability $1 - \alpha$, i.e. $\mathbb{P}(G_\alpha = k) = (1 - \alpha)\alpha^{k-1}$ for $k \geq 1$, independent of $\{Z_k^{rL}, k \geq 1\}$, and $m_\alpha \triangleq \lceil -\frac{1}{\log_2(\alpha)} \rceil$ denote a median of G_α . Note that the memoryless property implies $\mathbb{P}(G_\alpha \geq 2m_\alpha) \geq \frac{1}{4}$. Let $\xi_0 \triangleq 2^{-\frac{\hat{p}_0^2}{400}}$, and $\bar{\xi}_0 \triangleq 2^{-\frac{4}{\hat{p}_0^2 \eta_0^2} \epsilon_0^2}$.

LEMMA 9.

- (i) $.998 < \xi_0 \leq 1 - \epsilon_0 \leq \bar{\xi}_0 < 1$.
- (ii) $L \geq \epsilon_0^{-2}$ implies $\epsilon_0^{-3} L \exp(-\epsilon_0 L) \leq 25$.
- (iii) $\alpha \in [\xi_0, \bar{\xi}_0]$ implies:

- $m_\alpha, 2m_\alpha \in [400\hat{p}_0^{-2}, (\hat{p}_0 \eta_0 \epsilon_0^{-1})^2]$;

$$\bullet \quad \frac{1}{4}(1-\alpha)^{-1} \leq m_\alpha \leq 4(1-\alpha)^{-1}.$$

Proof of Lemma 9

i: That $\underline{\xi}_0 > .998$ follows from the fact that $\hat{p}_0 < 1$, and thus $2^{-\frac{\hat{p}_0^2}{400}} > 2^{-\frac{1}{400}} \geq .998$. That $\underline{\xi}_0 \leq 1 - \epsilon_0$ follows from the fact that (by definition) $\epsilon_0 \leq 1 - \underline{\xi}_0$. We now prove that $1 - \epsilon_0 \leq \bar{\xi}_0$. By definition $\epsilon_0 \leq \frac{1}{4}(\eta_0 \hat{p}_0)^2$, which implies that $\frac{4}{\hat{p}_0^2 \eta_0^2} \epsilon_0^2 \leq \epsilon_0$. Combining with the exponential inequality and the fact that $\log(2) < 1$, we conclude that

$$\bar{\xi}_0 \geq 2^{-\epsilon_0} \geq 1 - \log(2)\epsilon_0 \geq 1 - \epsilon_0,$$

completing the proof. As trivially $\bar{\xi}_0 < 1$, this completes the demonstration.

ii: It is easily verified that $\zeta_1(L) \triangleq L \exp(-\epsilon_0 L)$ is decreasing in L on $[\epsilon_0^{-1}, \infty)$. Thus $L \geq \epsilon_0^{-2}$ implies

$$\epsilon_0^{-3} L \exp(-\epsilon_0 L) \leq \epsilon_0^{-3} \zeta_1(\epsilon_0^{-2}) = \epsilon_0^{-5} \exp(-\epsilon_0^{-1}).$$

As $\zeta_2(\epsilon_0) \triangleq \epsilon_0^{-5} \exp(-\epsilon_0^{-1})$ is increasing in ϵ_0 on $(0, \frac{1}{5})$, and by definition $\epsilon_0 < \frac{1}{5}$, it follows that $\zeta_2(\epsilon_0) \leq \zeta_2(\frac{1}{5}) < 25$. Combining the above completes the proof.

iii: The first assertion follows immediately from the definitions of $\underline{\xi}_0, \bar{\xi}_0$, and m_α , and a straightforward calculation. To prove the second assertion, note that due to (i), $\alpha \in [\underline{\xi}_0, \bar{\xi}_0]$ implies $\alpha \in (.998, 1)$. It follows from a straightforward Taylor expansion of the logarithm function that $\alpha \in (.998, 1)$ implies $-2(1-\alpha) \leq \log_2(\alpha) \leq -(1-\alpha)$, and thus $\frac{1}{2}(1-\alpha)^{-1} \leq -\frac{1}{\log_2(\alpha)} \leq (1-\alpha)^{-1}$. Noting that $\alpha \in (.998, 1)$ implies $\lceil (1-\alpha)^{-1} \rceil \leq 4(1-\alpha)^{-1}$ completes the proof.

With Lemmas 8 and 9 in hand, we now complete the proof of Corollary 4.

Proof of Corollary 4 Suppose for contradiction that for some $L > \epsilon_0^{-2} + L_0 + 1$, it holds that $r_L > \mathbb{E}[D] - \epsilon_0$. In this case, it follows from Corollary 3, (1), and Jensen's inequality that for all $\alpha \in (\frac{1}{2}, 1)$,

$$\text{OPT}(L) \geq \min(b, h) \inf_{S \in \mathbb{R}} \sum_{k=1}^{\infty} (1-\alpha) \alpha^{k-1} |S + M_k^{rL}| - U_0 2^{L_0} (1-\alpha)^{-3} L \alpha^L. \quad (17)$$

Note that we may interpret the r.h.s. of (17) as an appropriate single-stage newsvendor problem (with ordering level S and demand distributed as $M_{G_\alpha}^{rL}$). We conclude from Lemmas 8 and 9, well-known results for the newsvendor problem (cf. Zipkin (2000)), and the memoryless property that for all $\alpha \in [\underline{\xi}_0, \bar{\xi}_0]$,

$$\begin{aligned}
\text{OPT}(L) &\geq \min(b, h) \mathbb{E}[|M_{m_\alpha}^{rL} - M_{G_\alpha}^{rL}|] - U_0 2^{L_0} (1 - \alpha)^{-3} L \alpha^L \\
&\geq \frac{1}{4} \min(b, h) (M_{2m_\alpha}^{rL} - M_{m_\alpha}^{rL}) - U_0 2^{L_0} (1 - \alpha)^{-3} L \alpha^L \\
&\geq \frac{1}{4} \min(b, h) \left(\frac{1}{5} \hat{p}_0 \eta_0 ((2m_\alpha)^{\frac{1}{2}} - m_\alpha^{\frac{1}{2}}) - \epsilon_0 m_\alpha - 2\eta_0 (\log(2) + 2) \right) - U_0 2^{L_0} (1 - \alpha)^{-3} L \alpha^L \\
&\geq \frac{1}{20} \min(b, h) \hat{p}_0 \eta_0 (2^{\frac{1}{2}} - 1) m_\alpha^{\frac{1}{2}} - \epsilon_0 m_\alpha - 6\eta_0 - U_0 2^{L_0} (1 - \alpha)^{-3} L \alpha^L \\
&\geq c_0 (1 - \alpha)^{-\frac{1}{2}} - 4\epsilon_0 (1 - \alpha)^{-1} - 6\eta_0 - U_0 2^{L_0} (1 - \alpha)^{-3} L \alpha^L.
\end{aligned}$$

Setting $\alpha = 1 - \epsilon_0$, and combining the above with Lemma 9.(i) and the fact that $1 - \epsilon_0 \leq \exp(-\epsilon_0)$, we conclude that

$$\text{OPT}(L) \geq c_0 \epsilon_0^{-\frac{1}{2}} - U_0 2^{L_0} \epsilon_0^{-3} L \exp(-\epsilon_0 L) - 6(\eta_0 + 1).$$

Applying Lemma 9.(ii) and the fact that $\text{OPT}(L) \leq U$, along with a straightforward contradiction argument, completes the proof. \square .